



US006115709A

United States Patent [19][11] **Patent Number:** **6,115,709****Gilmour et al.**[45] **Date of Patent:** **Sep. 5, 2000**

[54] **METHOD AND SYSTEM FOR CONSTRUCTING A KNOWLEDGE PROFILE OF A USER HAVING UNRESTRICTED AND RESTRICTED ACCESS PORTIONS ACCORDING TO RESPECTIVE LEVELS OF CONFIDENCE OF CONTENT OF THE PORTIONS**

[75] **Inventors:** David L. Gilmour, Los Altos Hills;
Hua-Wen Wang, Milpitas, both of
Calif.

[73] **Assignee:** Tacit Knowledge Systems, Inc., Palo
Alto, Calif.

[21] **Appl. No.:** 09/157,092

[22] **Filed:** Sep. 18, 1998

[51] **Int. Cl.⁷** G06F 17/30

[52] **U.S. Cl.** 707/9; 707/5; 707/6; 707/10;
705/7; 713/200

[58] **Field of Search** 707/9, 10, 1-2,
707/104; 709/217-219; 705/14, 26, 35-37;
713/200-202

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,914,586	4/1990	Swinehart et al.	364/200
5,247,575	9/1993	Sprague et al.	380/9
5,247,661	9/1993	Hager et al.	395/600
5,251,131	10/1993	Masand et al.	364/419.08

(List continued on next page.)

OTHER PUBLICATIONS

Welcome to enonymous.com, Web page, "Be Privacy Aware . . . Be enonymous", <http://www.enonymous.com/default.asp>. date unknown.

Enonymous Web Page, "Why be enonymous?", <http://www.enonymous.com/whybeenon.asp> date unknown.

Enonymous Web Page, What is enonymous advisor?, <http://www.enonymous.com/whatisit.asp>. date unknown.

Enonymous Web Page, "How enonymous advisor beta works . . .", <http://www.enonymous.com/howitworks.asp> date unknown.

Enonymous Web Page, "The enonymous zone . . .", <http://www.enonymous.com/zone.asp> date unknown.

Enonymous Web Page, "the enonymous profile . . .", <http://www.enonymous.com/profile.asp> date unknown.

Enonymous Web Page, "Frequently Asked Questions", <http://www.enonymous.com/faq.asp> date unknown.

"Yenta: A Multi-Agent, Referral-Based Matchmaking System", Leonard N. Foner, *The First International Conference on Autonomous Agents (Agents '97)*, Marina del Rey, CA, 1997.

"A Multi-Agent Referral System for Matchmaking", Leonard N. Foner, *The First International Conference on the Practical Applications of Intelligent Agents and Multi-Agent Technology*, London UK, Apr. 1996.

"Clustering and Information Sharing in an Ecology of Cooperating Agents", Leonard N. Foner, *AAAI Workshop on Information Gathering in Distributed, Heterogenous Environments '95*, Palo Alto, CA 1995.

(List continued on next page.)

Primary Examiner—Hosain T. Alam


Attorney, Agent, or Firm—Blakely, Sokoloff, Taylor & Zafman, LLP

[57]

ABSTRACT

A method of constructing a user knowledge profile, having distinct public and private portions with different access restrictions, requires assigning a confidence level to content within an electronic document. The electronic document is associated with a user, such as for example the author of the document. The content may be potentially indicative of a knowledge base of the user. The content is then stored in either the public or private portion of the user knowledge profile dependent upon whether the confidence level exceeds, or falls below, a predetermined threshold level. The public portion of the user knowledge profile is freely accessible by third parties, whereas the private portion has restricted access.

67 Claims, 32 Drawing Sheets



U.S. PATENT DOCUMENTS

5,276,869	1/1994	Forrest et al.	395/600
5,333,237	7/1994	Stefanopoulos et al.	395/12
5,473,732	12/1995	Chang	395/77
5,493,729	2/1996	Nigawara et al.	395/61
5,513,126	4/1996	Harkins et al.	364/514
5,541,836	7/1996	Church et al.	364/419.07
5,586,218	12/1996	Allen	395/10
5,608,900	3/1997	Dockter et al.	395/613
5,628,011	5/1997	Ahamed et al.	395/610
5,692,107	11/1997	Simoudis et al.	395/50
5,704,017	12/1997	Heckerman et al.	395/61
5,717,914	2/1998	Husick et al.	395/605
5,717,923	2/1998	Dedrick	395/613
5,720,001	2/1998	Nguyen	395/10
5,724,567	3/1998	Rose et al.	707/2
5,727,129	3/1998	Barrett et al.	395/12
5,754,938	5/1998	Herz et al.	455/4.2
5,794,210	8/1998	Goldhaber et al.	705/14
5,802,518	9/1998	Karaev et al.	707/9
5,855,008	12/1998	Goldhaber et al.	705/14
5,913,212	6/1999	Sutcliffe et al.	707/6
5,931,907	8/1999	Davis et al.	709/218
5,999,975	12/1999	Kittaka et al.	709/224
6,021,439	2/2000	Turek et al.	709/224

OTHER PUBLICATIONS

"Somewhat-by-topic linearization of Yenta", Leonard Foner, <http://www.media.mit.edu/people/foner/Yenta/linearization-by-topic.html>.

"Political Artifacts and Personal Privacy: The Yenta Multi-Agent Distributed Matchmaking System", Leonard Newton Foner, Apr. 30, 1999, © Massachusetts Institute of Technology, 1999.

"IS Puts Notes to the Test", *Datamation*, Mark Schlack, vol. 37, No. 15, pp 24-26, Aug. 1, 1991.

"Mail-Man: A Knowledge-Based Mail Assistant for Managers", *Journal of Organizational Computing*, L.F. Motiwalla and J.F. Nunamaker, Jr., vol. 2, No. 2, pp 131-154, 1992, (abstract only).

"Topic Real-Time", HP-UX-Documents Disc 50726-10186 (from Software Patent Institute Database of Software Technologies), Feb. 1, 1995.

"askSam for Windows Getting Started Guide", (from Software Patent Institute Database of Software Technologies), first section, May 30, 1995.

"EZ Reader: Embedded AI for Automatic Electronic Mail Interpretation and Routing", *Proceedings of the Thirteenth National Conference on Artificial Intelligence and the Eighth Innovative Applications of Artificial Intelligence Conference*, A. Rice, J. Hsu, A. Angotti and R. Piccolo, vol. 2, pp 1507-1517, Aug. 4-8, 1996 (INSPEC Abstract).

"Information Management for Knowledge Amplification in Virtual Enterprises", J. Numata et al., *IEMC Proceedings*, pp 281-285, Aug. 18-20, 1996 (INSPEC Abstract).

Data Sources. PP K-213, 214, 217, 224, 229 and 230, 1st Edition 1998.

<http://www.email-software.com/pages/00108.htm>, review of "Emailrobot for Exchange/SMTP".

<http://www.email-software.com/pages/00033.htm>, review of "Signup V2.0".

"Knowledge Management: Fuel For Innovation", Bob Evans, CPM net Information Week On Line, Oct. 20, 1997.

"Knowing What We Know", Justin Hibbard, CPM net Information Week On Line, Oct. 20, 1997.

"Knowledge Management Evaluation Scenario", Jeff Angus with Jeetu Patel and Joe Fenner, CPM net Information Week On Line, Mar. 16, 1998.

"Knowledge Management Takes Industry's Center Stage", Elliot Maise, CMP net Computer Reseller News, Feb. 2, 1998, Issue 774.

"Knowledge Management's Net Gain", Kevin Jones, ZD net, Inter @ctive Week, Feb. 24, 1998.

"Knowledge Management: Great Concept . . . But What Is It?", Jeff Angus, Jeetu Patel and Jennifer Harty, CMP net, Information Week On Line, Mar. 16, 1998.

"Open Sesame Site Just Works", Bil Burke, BusinessToday.com, <http://www.opensesame.com>, Today's Column, Jan. 22, 1998.

"New eCommerce and Entertainment Web Site Demonstrates Leading Edge Personalization and Privacy Features" (press release), <http://www.opensesame.com>, Jan. 20, 1998.

"Neural Agent Enables Personalized Surfing", R. Colin Johnson, CMP net TechWeb, <http://www.techweb.com>, Feb. 4, 1998.

"Open Sesame and Verity Open Doors to Personalized Software", KMWorld.com, <http://www.kmworld.com>, Feb. 6, 1998.

"Life Span vs Life Spam", George Gilder, Forbes ASAP, <http://www.forbes.com/asap>, Apr. 6, 1998.

"Learn Sesame gets more personal", Jim Rapoza, PC Week Online, <http://www.zdnet.com/pcweek/reviews>, Mar. 18, 1998.

"Natrificial Software Technologies Unveils Internet Brain Publishing" (press release), Natrificial Software Technologies, <http://www.natrificial.com>, Jun. 9, 1998.

"Natrificial Software Technologies Introduces The Brain" (press release), Natrificial Software Technologies, <http://www.natrificial.com>, Jan. 26, 1998.

"The Brain: Much More The Way You Think", Scot Finnie, CMP Net Windows Magazine, <http://www.windowsmagazine.com>, May 1, 1998.

"Your Brain on Windows", Leslie Ayers, ZD Net Products, <http://www.zdnet.com/products>, Apr., 1998.

"What's New", Natrificial Software Technologies, <http://www.natrificial.com>, 1998.

"Natrificial Software Technologies", Natrificial Software Technologies, <http://www.natrificial.com>.

"Always Thinking Ahead", Natrificial Software Technologies, <http://www.natrificial.com>.

"Digitize Your Mind", Natrificial Software Technologies, <http://www.natrificial.com>.

"It's Your Thought That Counts", Natrificial Software Technologies, <http://www.natrificial.com>.

"Knowledge if Power", Natrificial Software Technologies, <http://www.natrificial.com>.

"The Web at the Speed of Thought", Natrificial Software Technologies, <http://www.natrificial.com>.

"Free Your Mind", Natrificial Software Technologies, <http://www.natrificial.com>.

"The Brain Press Reviews", Natrificial Software Technologies, <http://www.natrificial.com>.

"Abuzz's Mission", Abuzz, <http://www.abuzz.com>.

"Beehive", Abuzz, <http://www.abuzz.com>.

"Beehive", Abuzz, <http://www.abuzz.com/home/benefits.htm>.

"Beehive: Beta Program", Abuzz, <http://www.abuzz.com/home/demos.htm>.

"Beehive: White Papers", Abuzz, http://www.abuzz.com/home/white_papers.htm.

Open Sesame, <http://www.opensesame.com>.
 Open Sesame: The Company, <http://www.opensesame.com/company.html>.
 Open Sesame: Published Papers, http://www.opensesame.com/co_02.html.
 Open Sesame: Agent Sourcebook, http://www.opensesame.com/co_03.html.
 Open Sesame: Products; Learn Sesame, <http://www.opensesame.com/products.html>.
 Open Sesame: Products; Why Personalize?, http://www.opensesame.com/prod_01.html.
 Open Sesame: Products: Benefits, http://www.opensesame.com/prod_02.html.
 Open Sesame: Products; Competitive Summary, http://www.opensesame.com/prod_03.html.
 Open Sesame: Products; Product Datasheet, http://www.opensesame.com/prod_04.html.
 Open Sesame: Products/Demos, http://www.opensesame.com/product_05.html.
 e-Genie by Open Sesame: Your Entertainment Genie!, <http://egenie.opensesame.com/>.
 Open Sesame: Products; eGenie Live!, http://www.opensesame.com/prod_06.html.
 Open Sesame: Commitment to Privacy, <http://www.opensesame.com/privacy.html>.
 Open Sesame: Press, <http://opensesame.com/press.cfm>.
 Open Sesame: Press; Reviews, http://www.opensesame.com/pr_02.html.
 Autonomy Knowledge Management Products, <http://www.agentware.com>.
 Autonomy Latest News, Autonomy Home Page, <http://www.agentware.com>.
 Autonomy Knowledge Server Data Sheet, <http://www.agentware.com>.
 Autonomy Knowledge Update Data Sheet, <http://www.agentware.com>.
 The Technology Behind Autonomy Agentware, Autonomy Technology Page, <http://www.agentware.com>.
 Autonomy Agentware Technology White Paper, <http://www.agentware.com>.
 "Autonomy Launches New Knowledge Management Products to Help Companies Leverage Employee Expertise, Late-Breaking News and Existing Information Archives", Autonomy. Press Release, Feb. 17, 1998, <http://www.agentware.com>.
 "Autonomy Ships Agentware Products for Knowledge Management", Autonomy Press Release, Apr. 14, 1998, <http://www.agentware.com>.
 "A Short Introduction to NPtool", Atro Voutilainen, A Short Introduction to NPtool, <http://www.lingsoft.fi/doc/nptool/intro/>.
 NPtool Intro: Overview, <http://www.lingsoft.fi/doc/nptool/intro/overview.html>.
 NPtool Intro: Use of a Noun Phrase Parser, <http://www.lingsoft.fi/doc/nptool/intro/uses.html>.
 NPtool Intro: Previous Work, <http://www.lingsoft.fi/doc/nptool/intro/previous.html>.
 NPtool Intro: NPtool in Outline, <http://www.lingsoft.fi/doc/nptool/intro/outline.html>.
 NPtool Intro: Syntactic Description, <http://www.lingsoft.fi/doc/nptool/intro/syntax.html>.
 NPtool Intro: Performance, <http://www.lingsoft.fi/doc/nptool/intro/performance.html>.
 NPtool Intro: Conclusion and Acknowledgements, <http://www.lingsoft.fi/doc/nptool/intro/conclusion.html>.
 Orbital Technologies: Our Focus, <http://www.orbital-tech.com>.
 Orbital Technologies: About Orbital Technologies, <http://www.orbital-tech.com/about.html>.
 Orbital Technologies: Products, <http://www.orbital-tech.com/products.html>.
 Orbital Technologies: Organik KnowledgeWare, <http://www.orbital-tech.com/organikkw.html>.
 Orbital Technologies: Organik Persona Server, <http://www.orbital-tech.com/organikps.html>.
 Orbital Technologies: Corporate Fact Sheet, <http://www.orbital-tech.com>.
 Orbital Technologies: Orbital Organik White Paper, Jul., 1997, <http://www.orbital-tech.com>.




FIG. 1

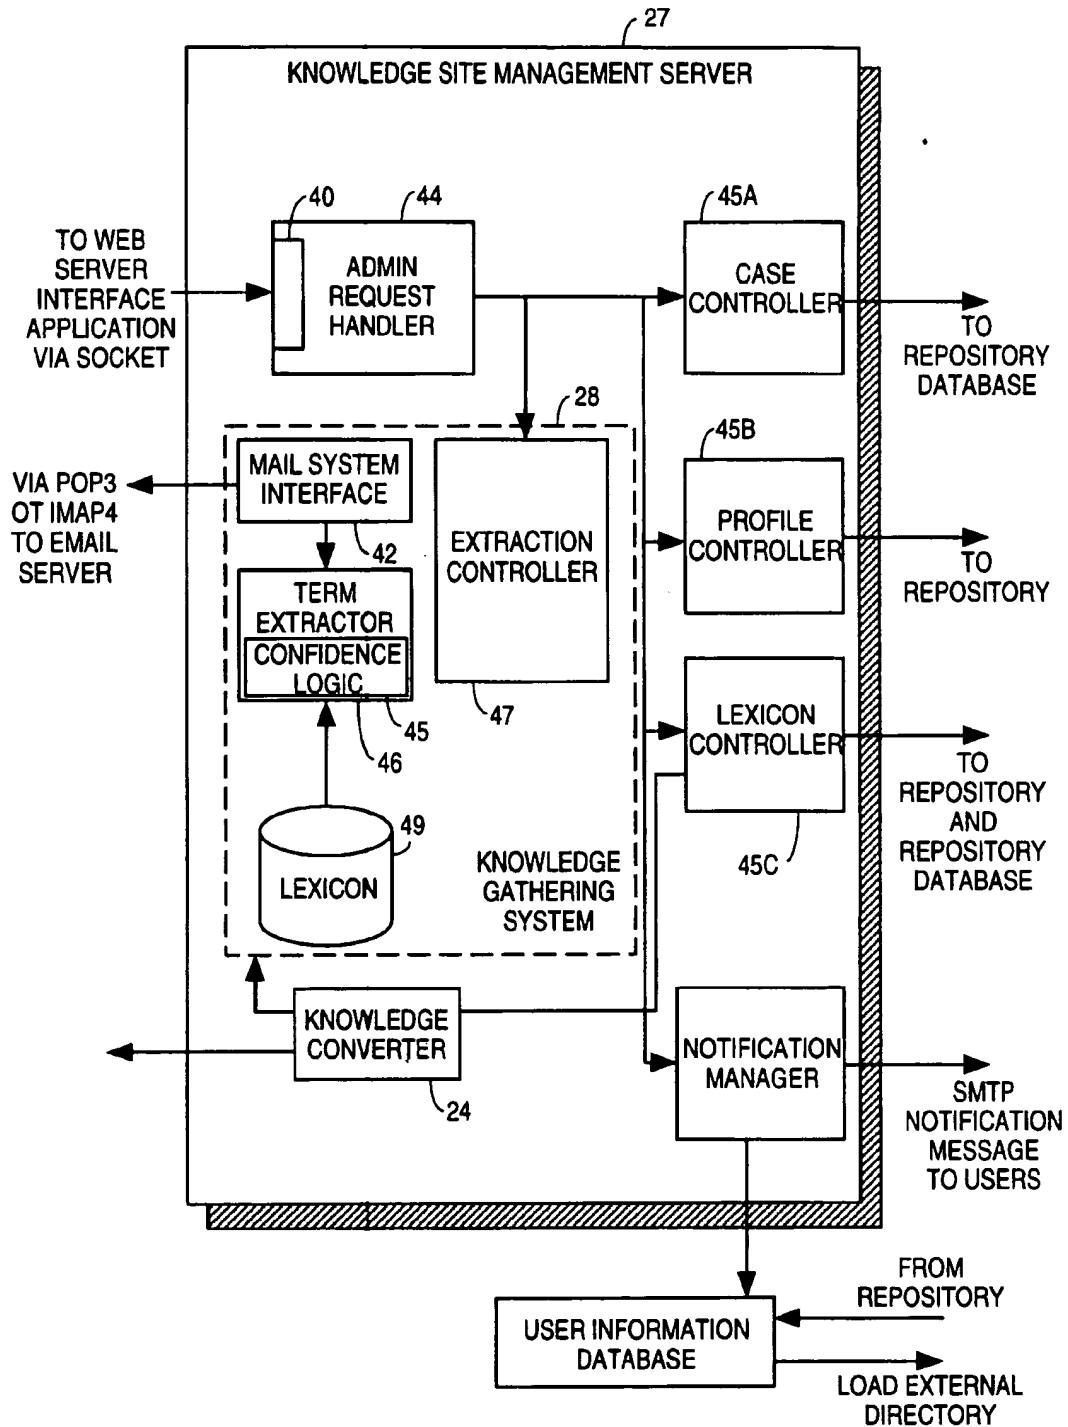





FIG. 2



**FIG. 5**




FIG. 6

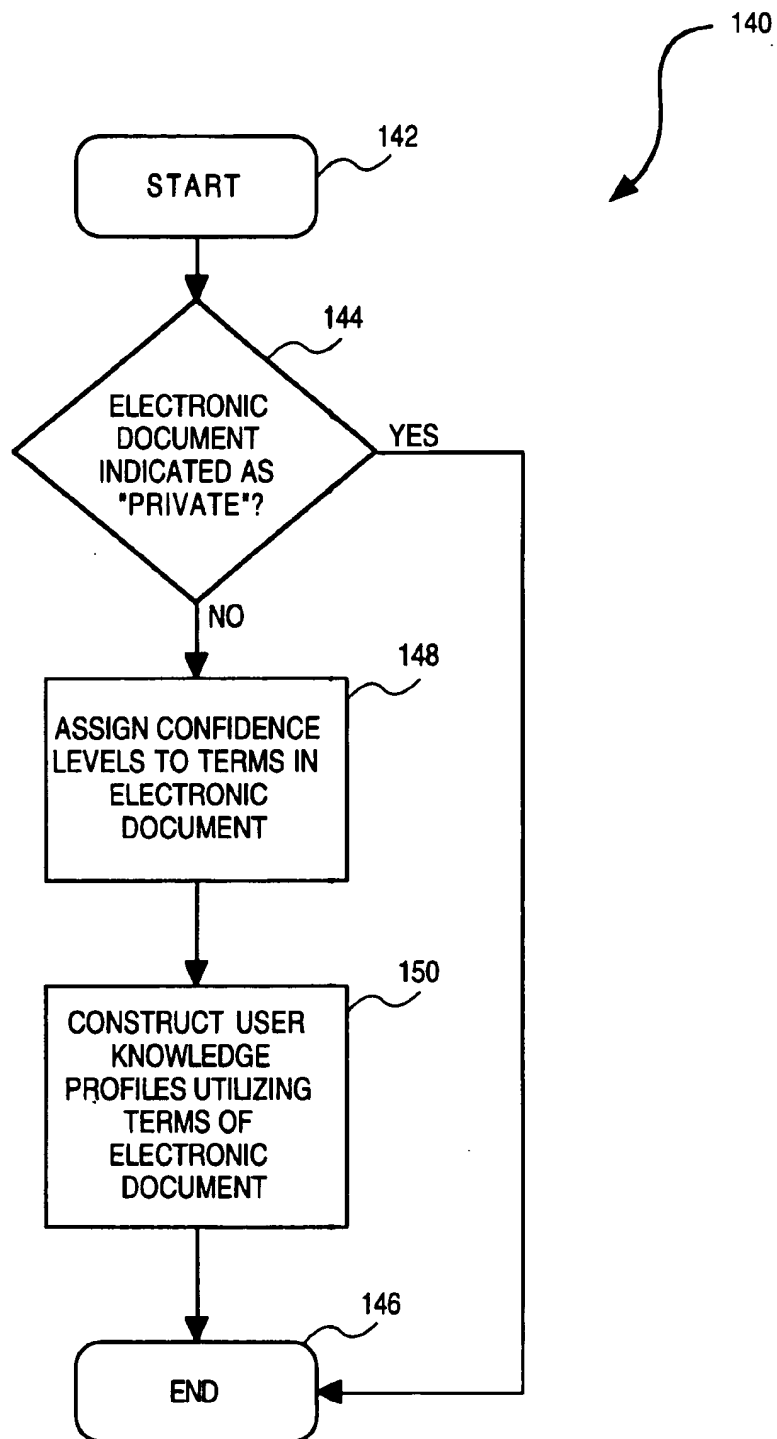


FIG. 7

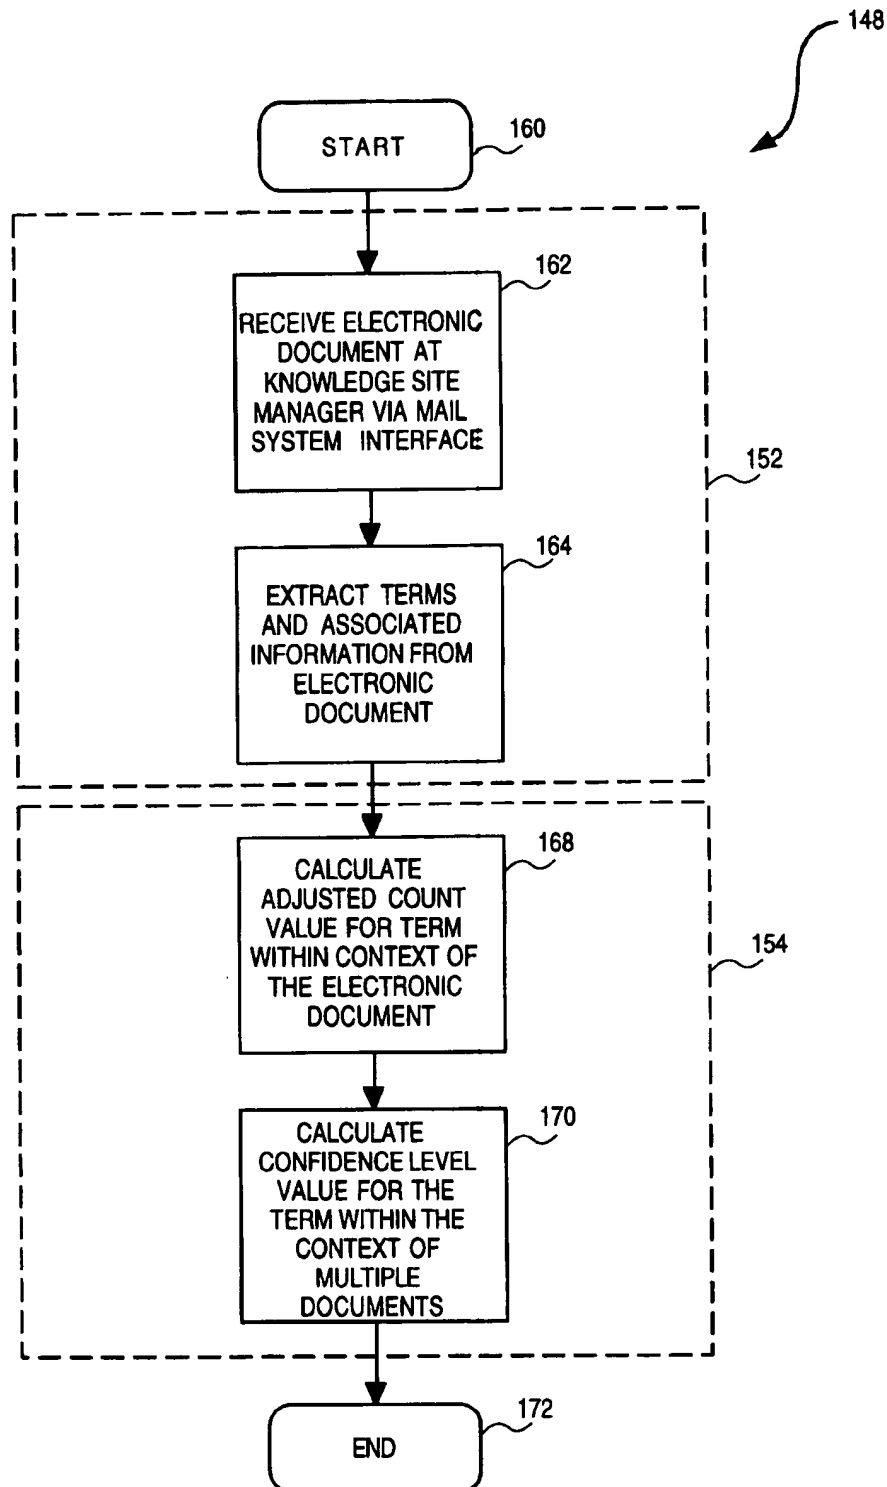


FIG. 8

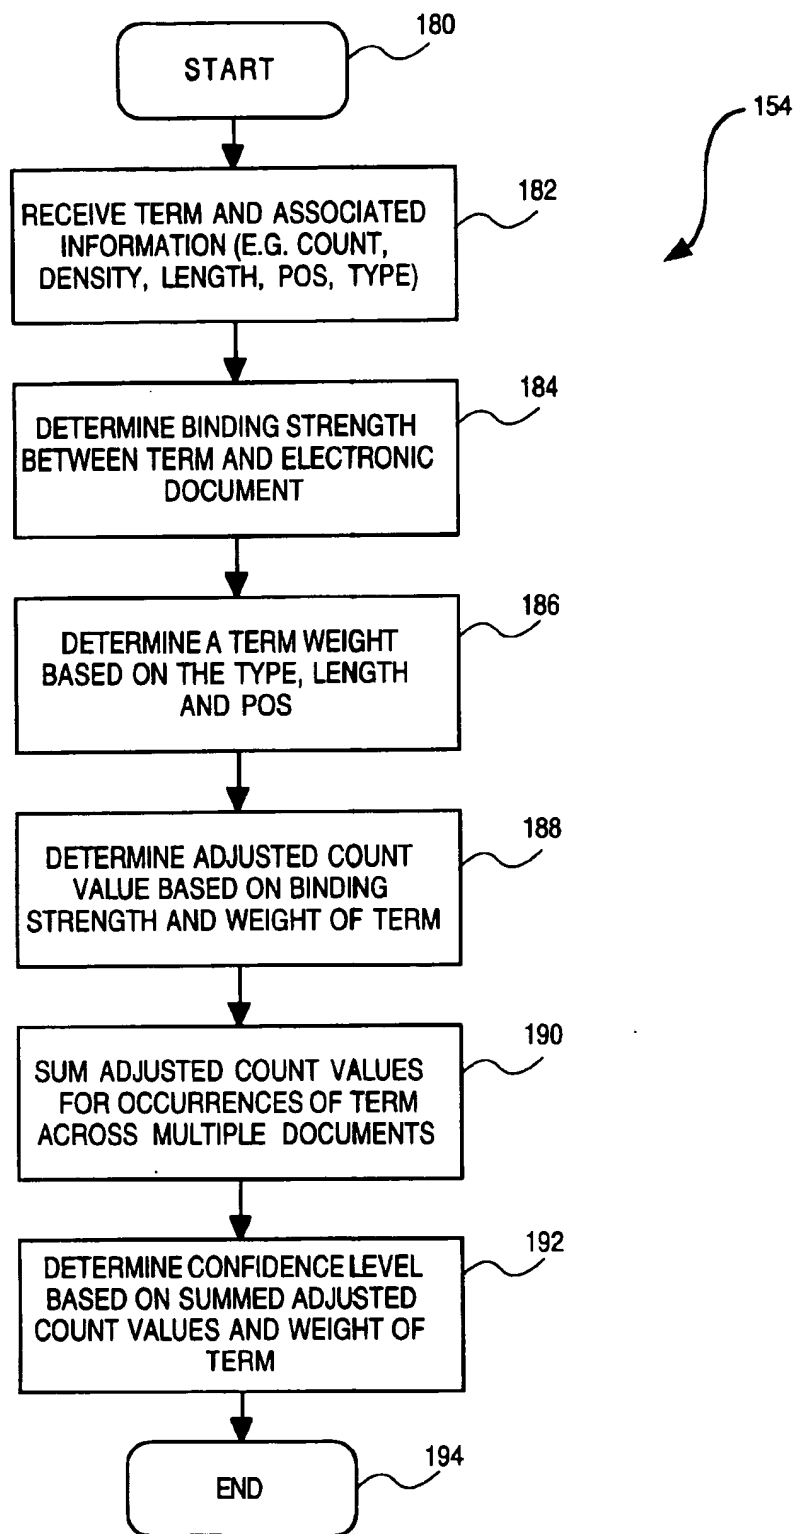


FIG. 9A




FIG. 9B

200

TERM-DOCUMENT BINDING TABLE		
CLASS	NO. OF OCCURRENCES	DENSITY
A	X	4.00%
B	X	2.00%
C	X	1.00%
D	X	0.5%
E	3	X
F	1	X

206 202 204

FIG. 10

210

WEIGHT TABLE				
WEIGHT	TYPE	LENGTH	POS	OTHER
10	P	1	A	
30	L	1	V	
150	P	2	NN	
300	L	2	N	
400	L	3	N	
450	U	7	X	

212 214 216 218

FIG. 11


220



	A	B	C	D	E	F
10	0	0	0	0	0	0
30	0	0	0	0	0	0
100	0.4	0.3	0.2	0.1	0.05	0
150	0.8	0.7	0.6	0.5	0.4	0.1
300	1.5	1.4	1.3	1	0.5	0.1
320	1.6	1.5	1.4	1.1	0.6	0.1
400	1.9	1.8	1.5	1.2	0.7	0.1
450	2	1.9	1.6	1.3	1	0.1

FIG. 12

230



	0.063	0.125	>>>>	32	64	128
10	0.1	0.2	>>>>	55	100	200
30	0.3	0.6	>>>>	180	350	600
100	1.6	2.8	>>>>	600	1000	2100
150	2.5	5	>>>>	1100	2000	5000
300	6.5	11	>>>>	2700	6000	15000
320	7.1	11.5	>>>>	2900	7000	16000
400	8	15	>>>>	3500	8000	20000
450	9	18	>>>>	4200	9000	24000

FIG. 13

240



	0.063	0.125	>>>>>	32	64	128
10	0	0	>>>>>	50	100	200
30	0	0	>>>>>	150	350	600
100	0	0	>>>>>	600	1000	1000
150	0	0	>>>>>	1000	1000	1000
300	0	0	>>>>>	1000	1000	1000
320	0	0	>>>>>	1000	1000	1000
400	0	25	>>>>>	1000	1000	1000
450	0	25	>>>>>	1000	1000	1000

FIG. 14

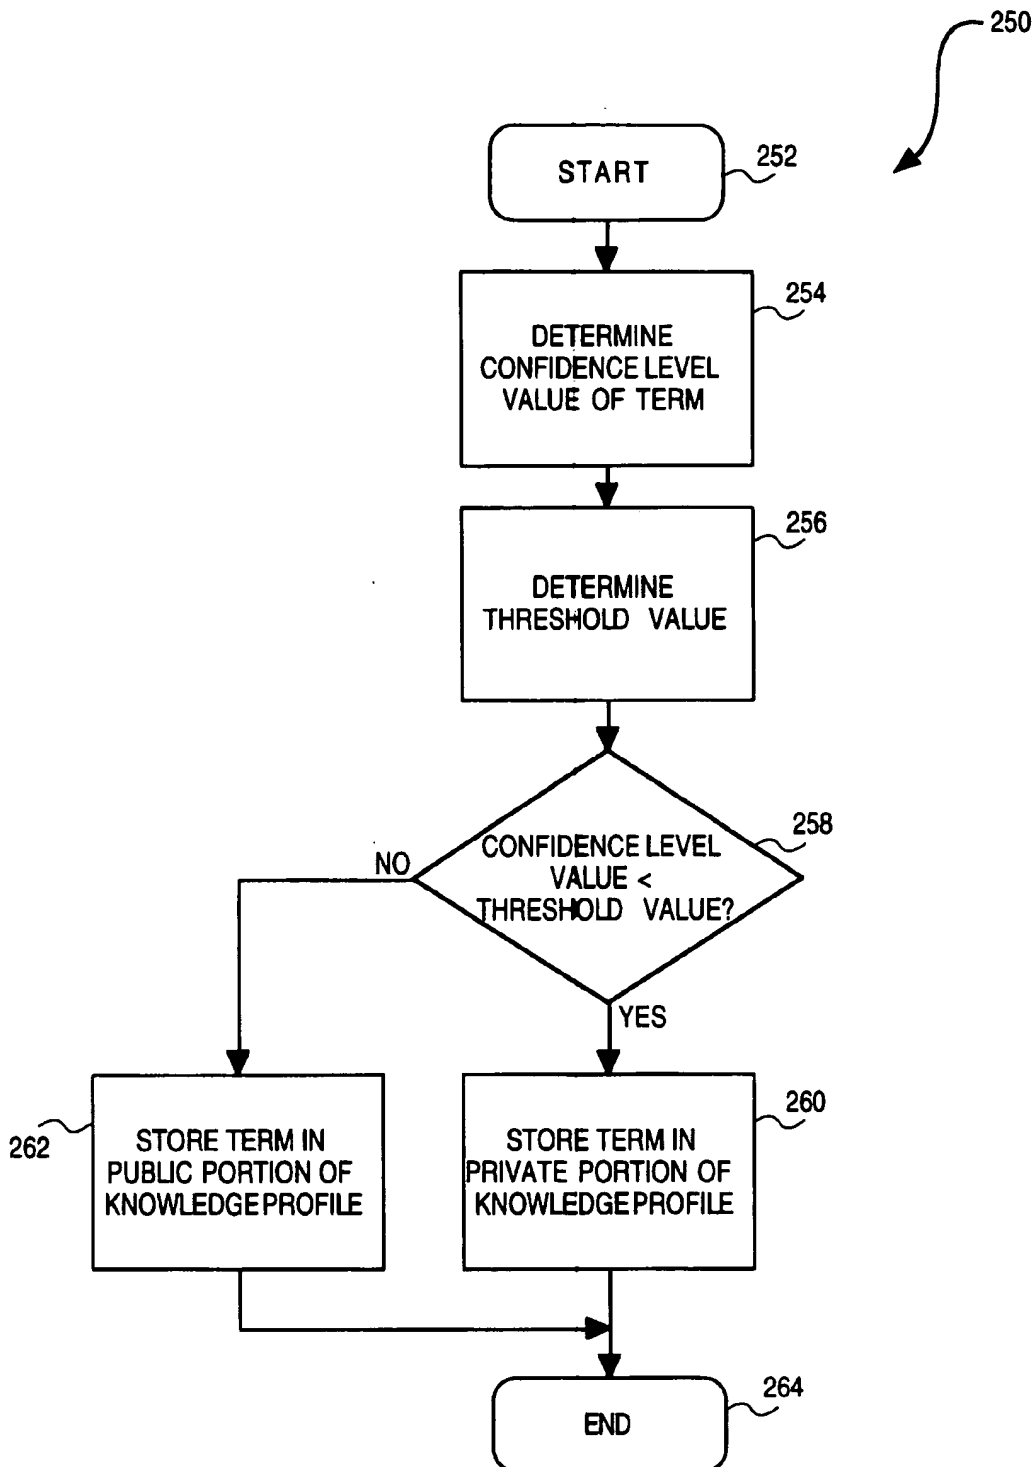
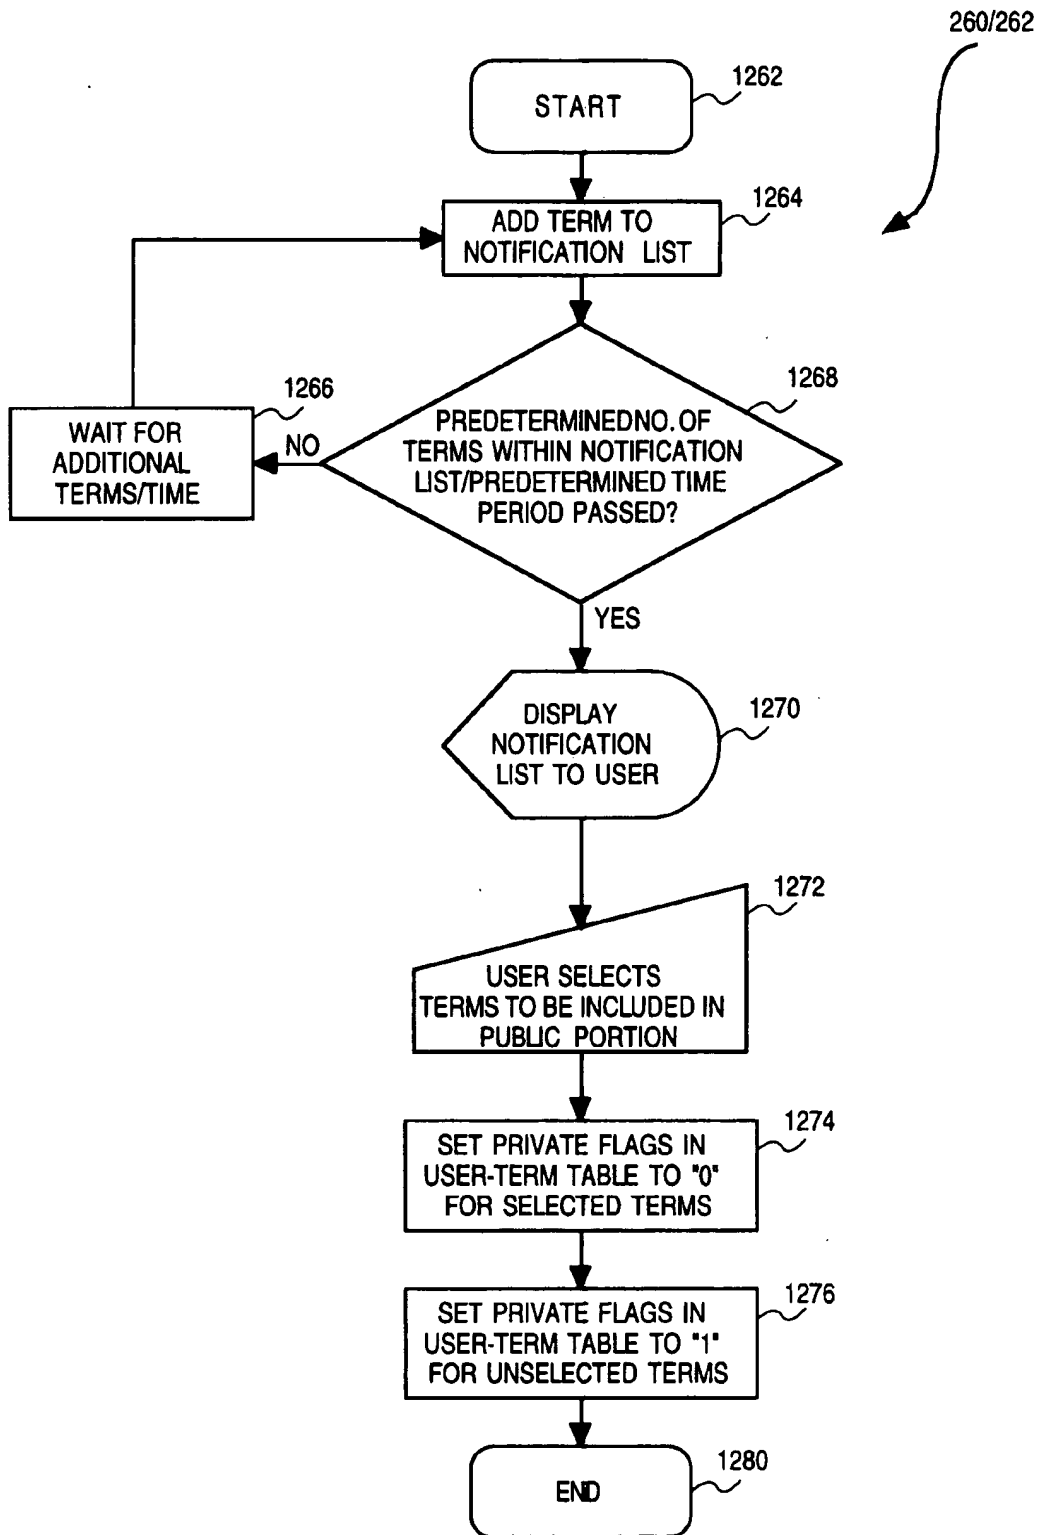


FIG. 15A


**FIG. 15B**

112

USER ID	TERM ID	CONFIDENCE LEVEL
1	101	990
1	106	976
1	207	950
.	.	.
.	.	.
.	.	.
1	67	310
1	98	300
1	116	290
1	190	287
.	.	.
.	.	.
.	.	.
8	107	995
8	128	992
8	154	900
.	.	.
.	.	.
.	.	.
8	62	370
8	8	300
8	52	207
8	261	200
.	.	.
.	.	.
.	.	.

92 94 118

FIG. 16A



USER ID	TERM ID	CONFIDENCE LEVEL	PRIVATE FLAG
1	101	990	0
1	106	976	0
1	207	950	1
.	.	.	.
.	.	.	.
.	.	.	.
1	67	310	0
1	98	300	1
1	116	290	0
1	190	287	0
.	.	.	.
.	.	.	.
.	.	.	.
8	107	995	1
8	128	992	0
8	154	900	0
.	.	.	.
.	.	.	.
.	.	.	.
8	62	370	0
8	8	300	0
8	52	207	0
8	261	200	0
.	.	.	.
.	.	.	.
.	.	.	.

92 94 118 119

FIG. 16B

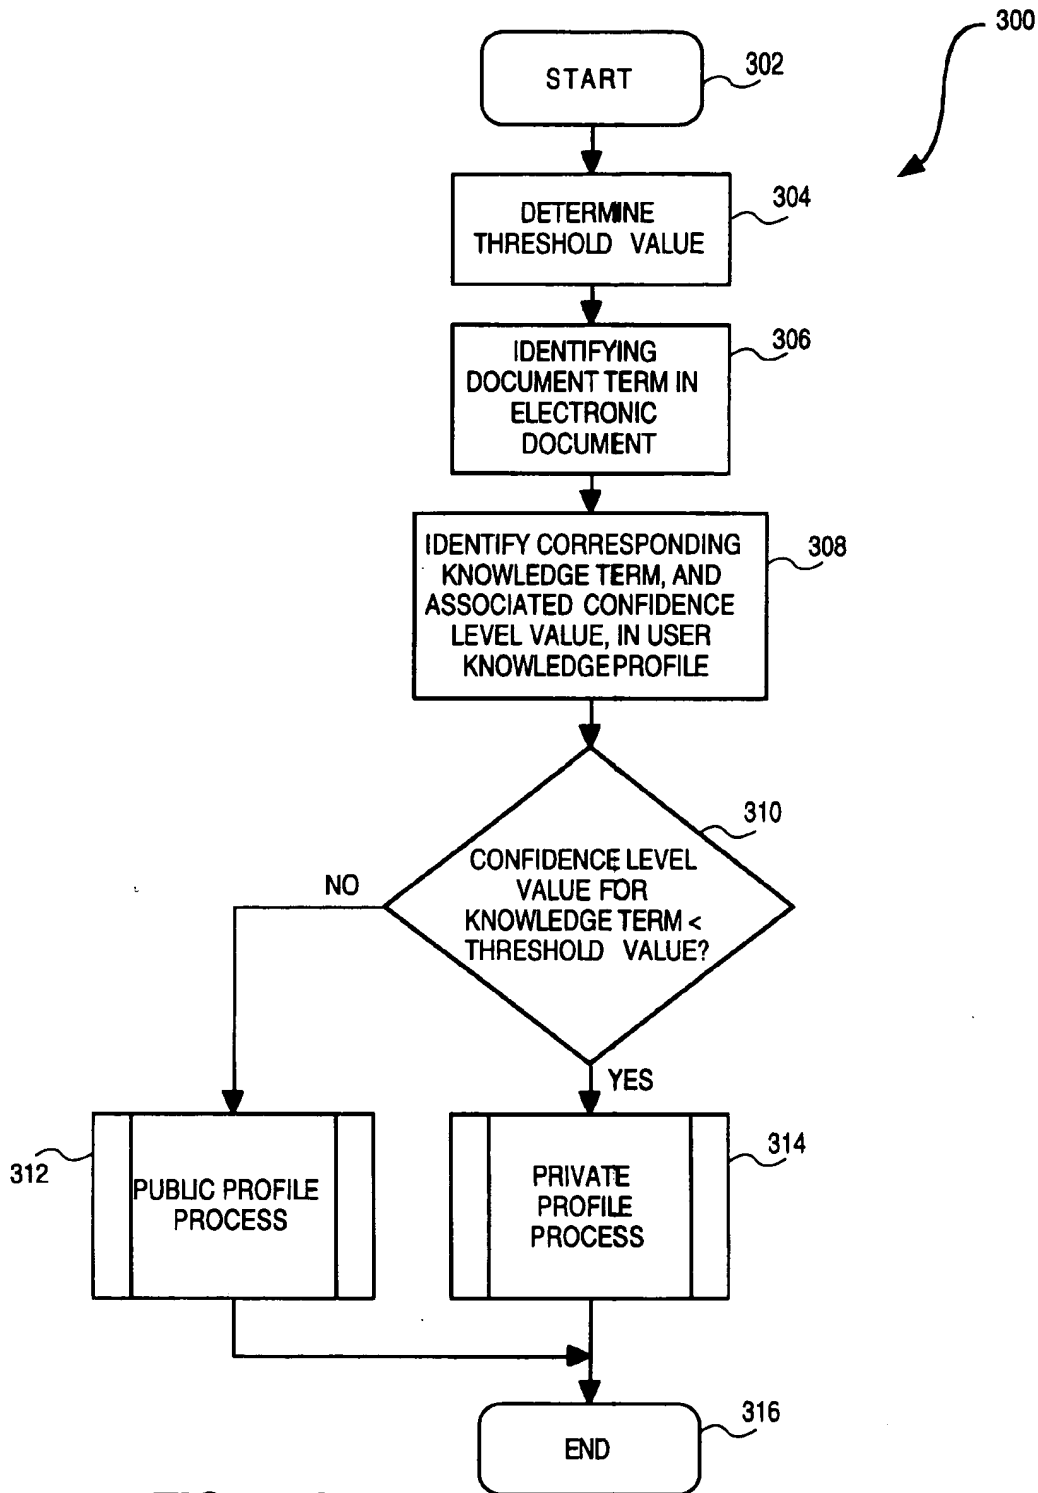
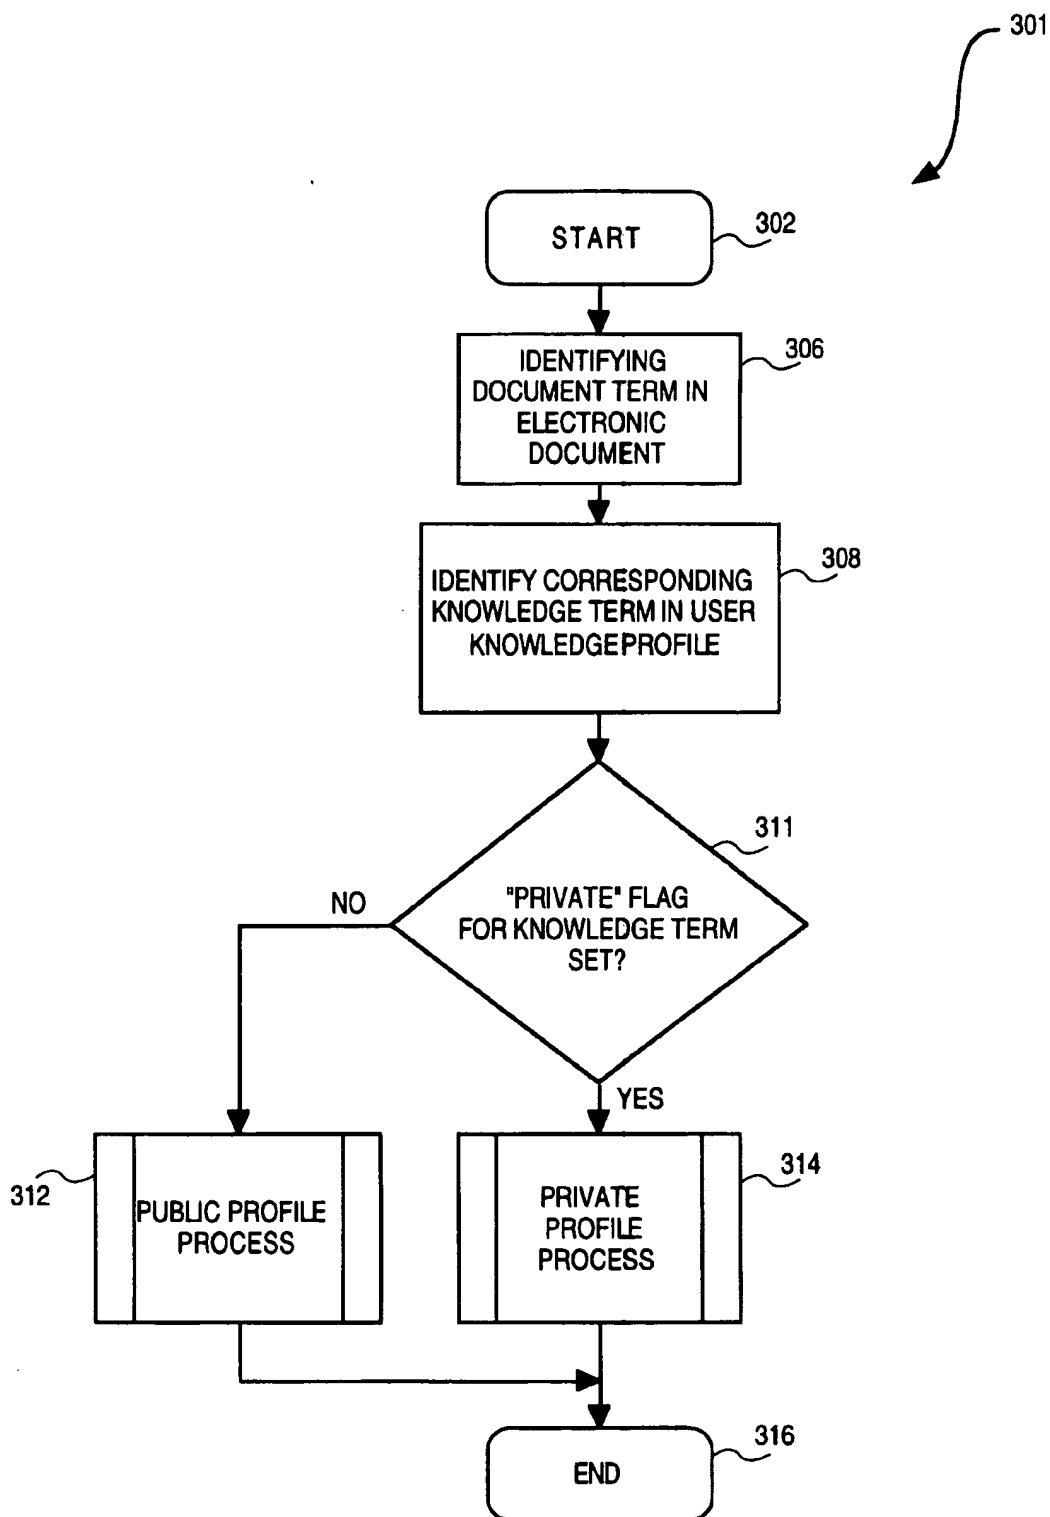
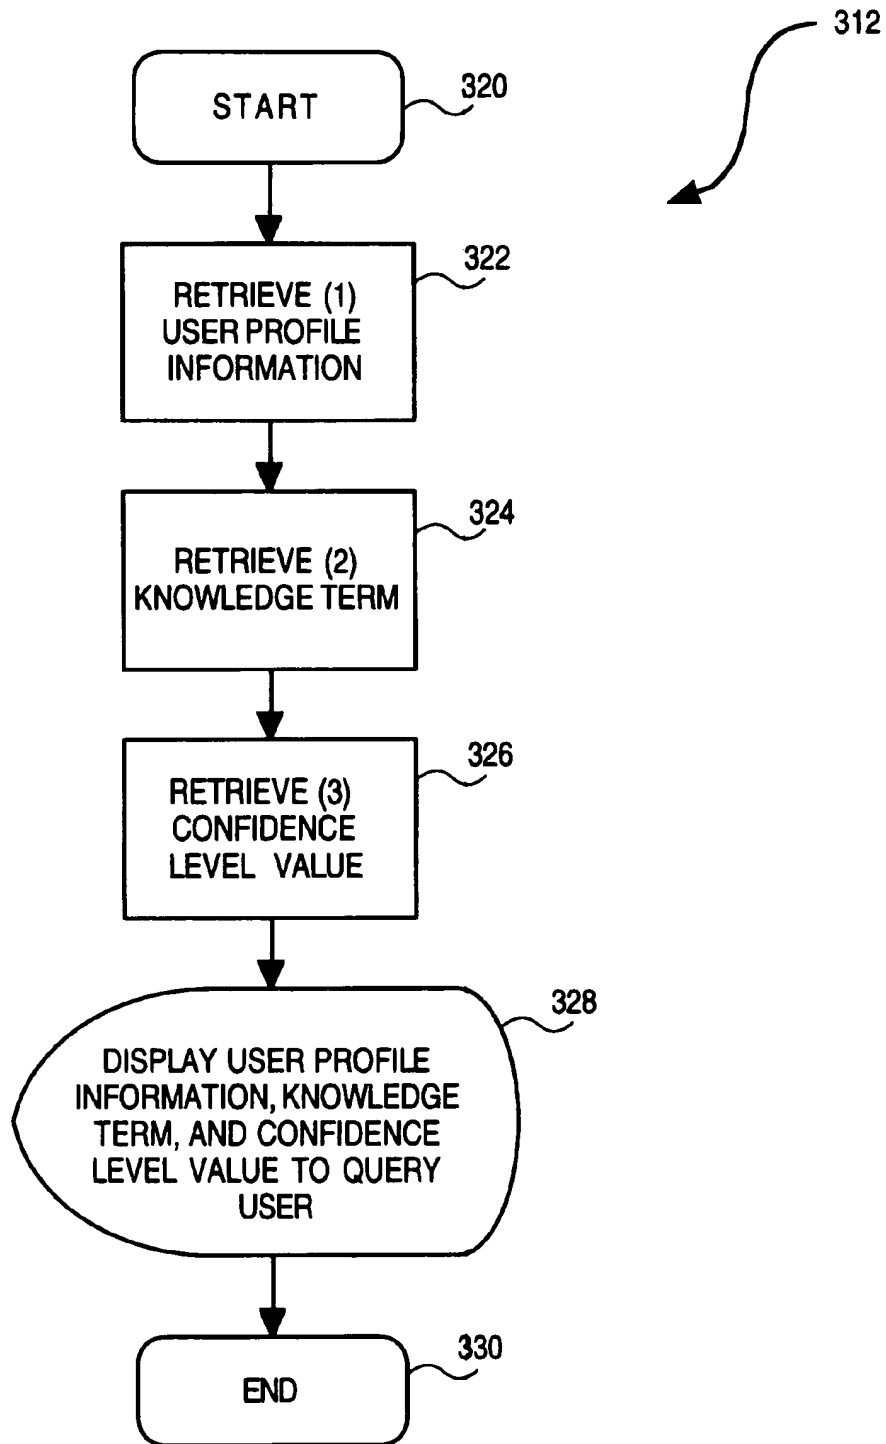
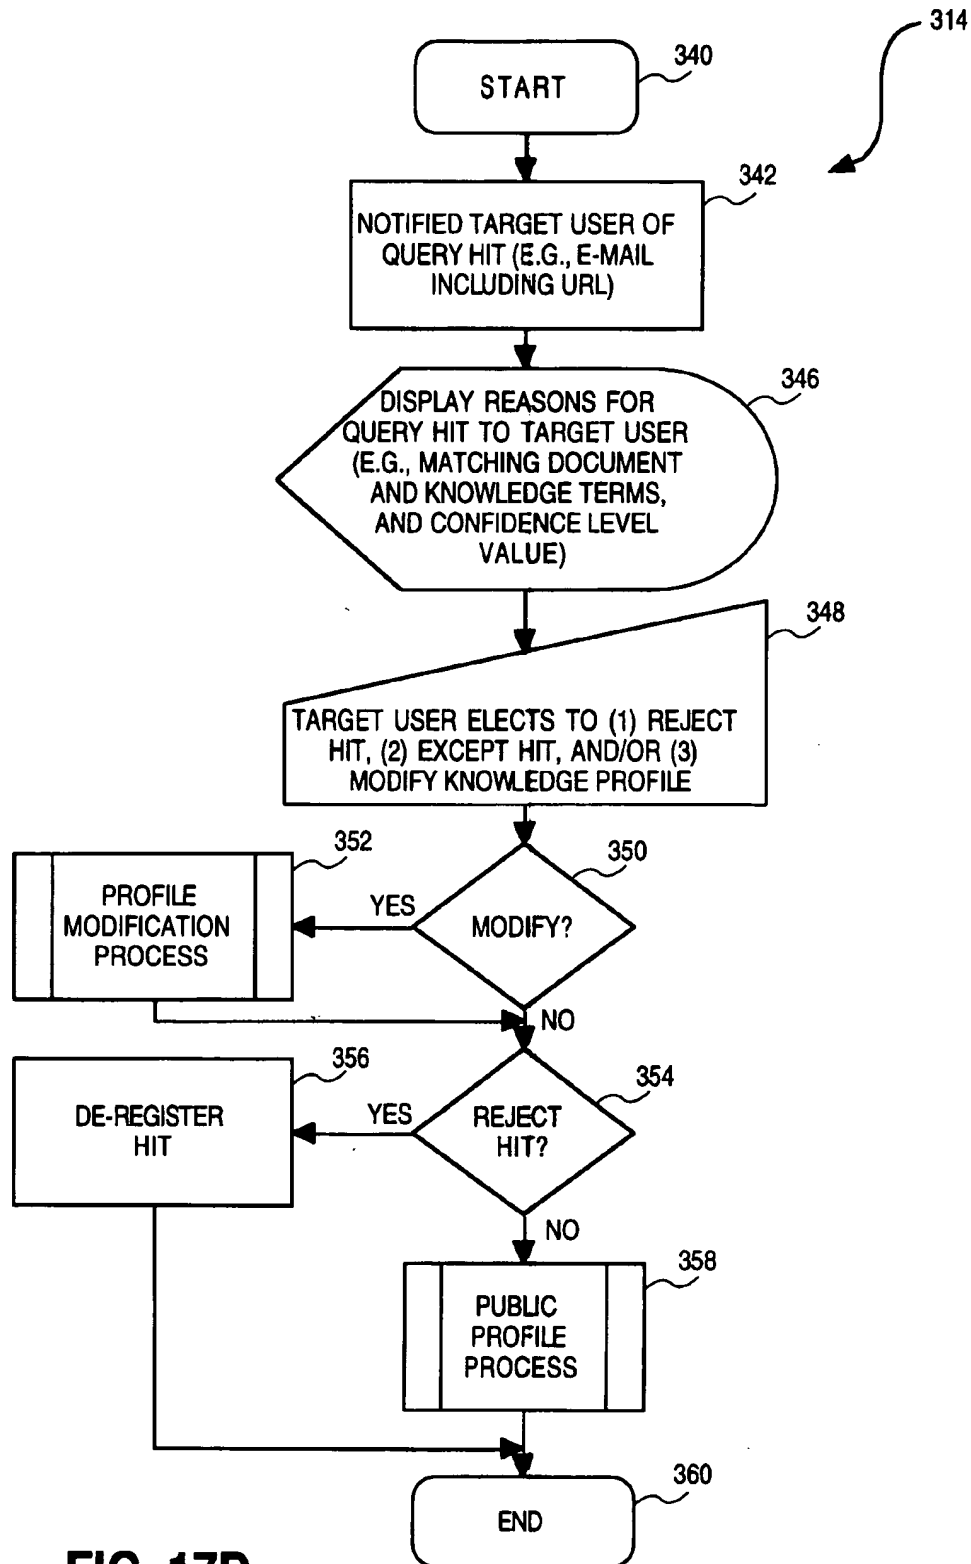


FIG. 17A

**FIG. 17B**

**FIG. 17C**



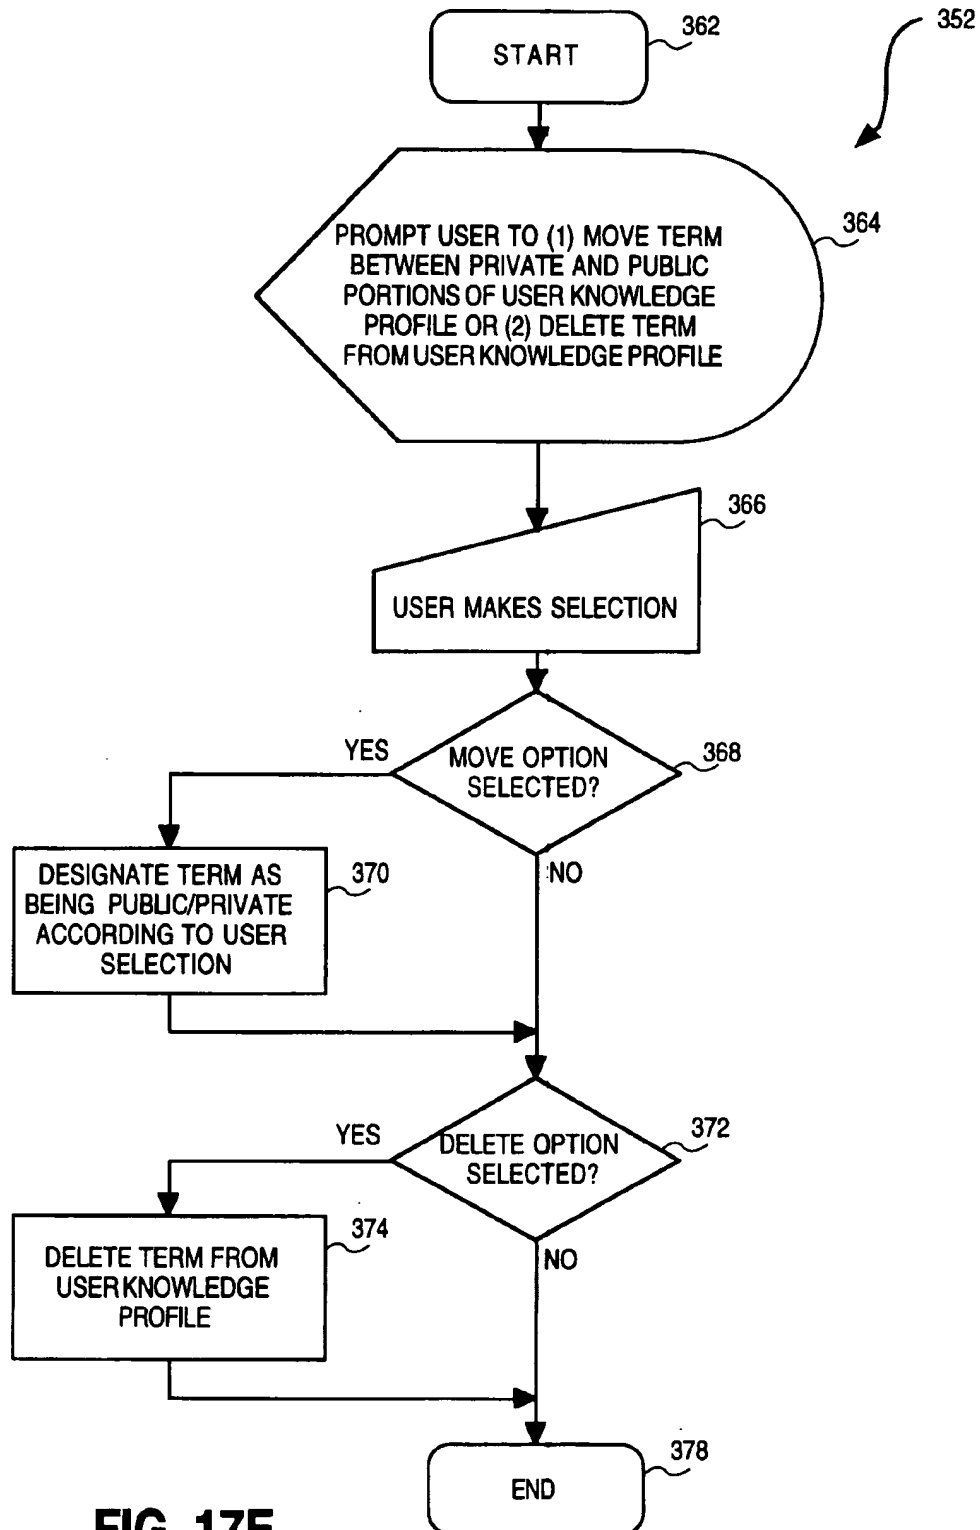
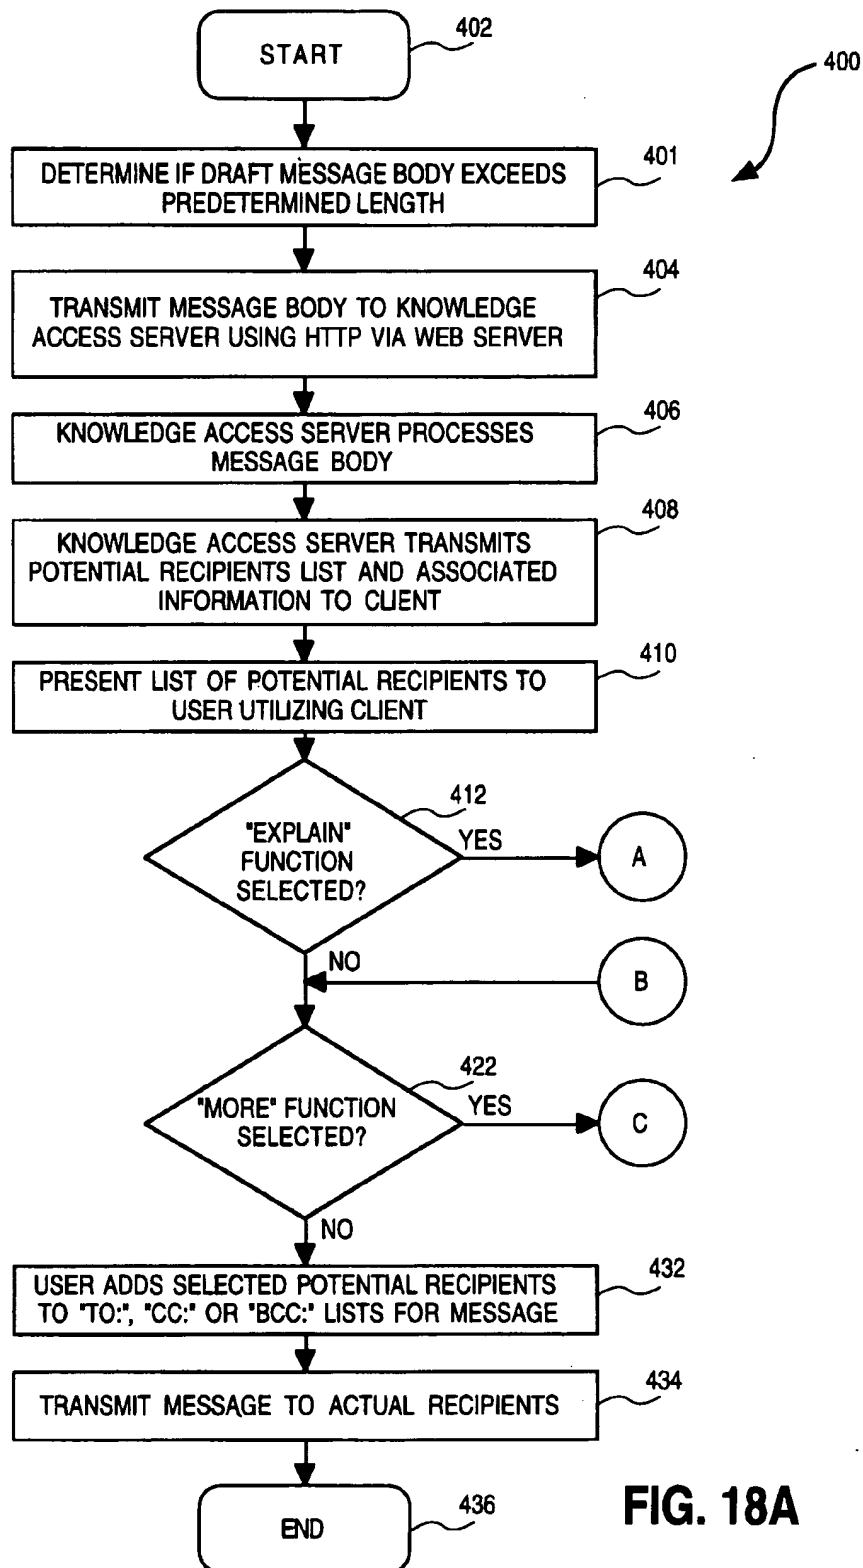
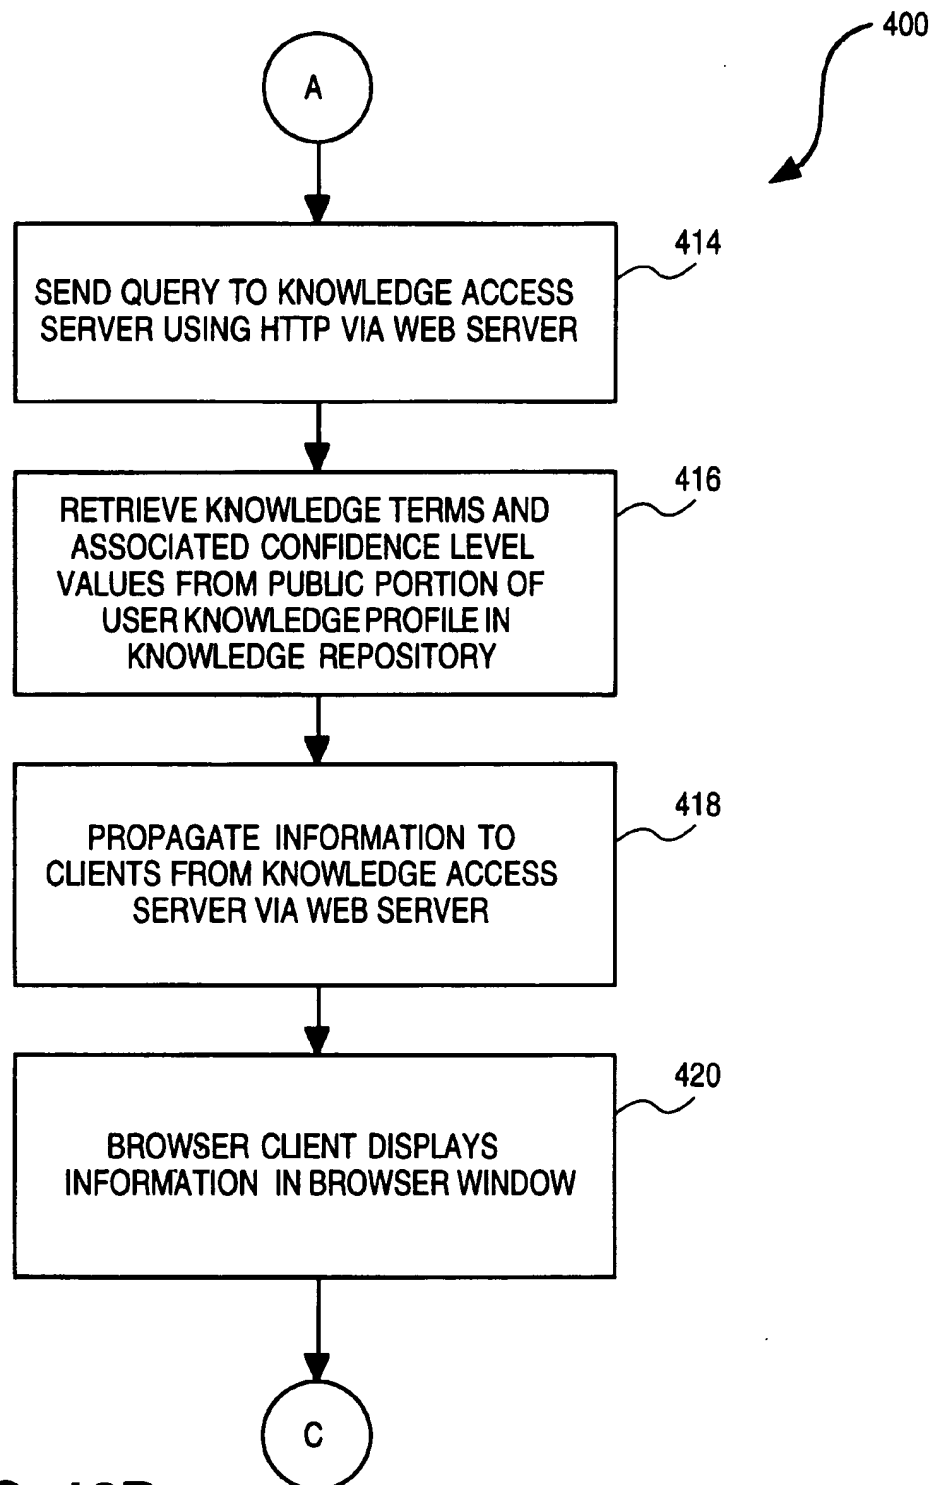
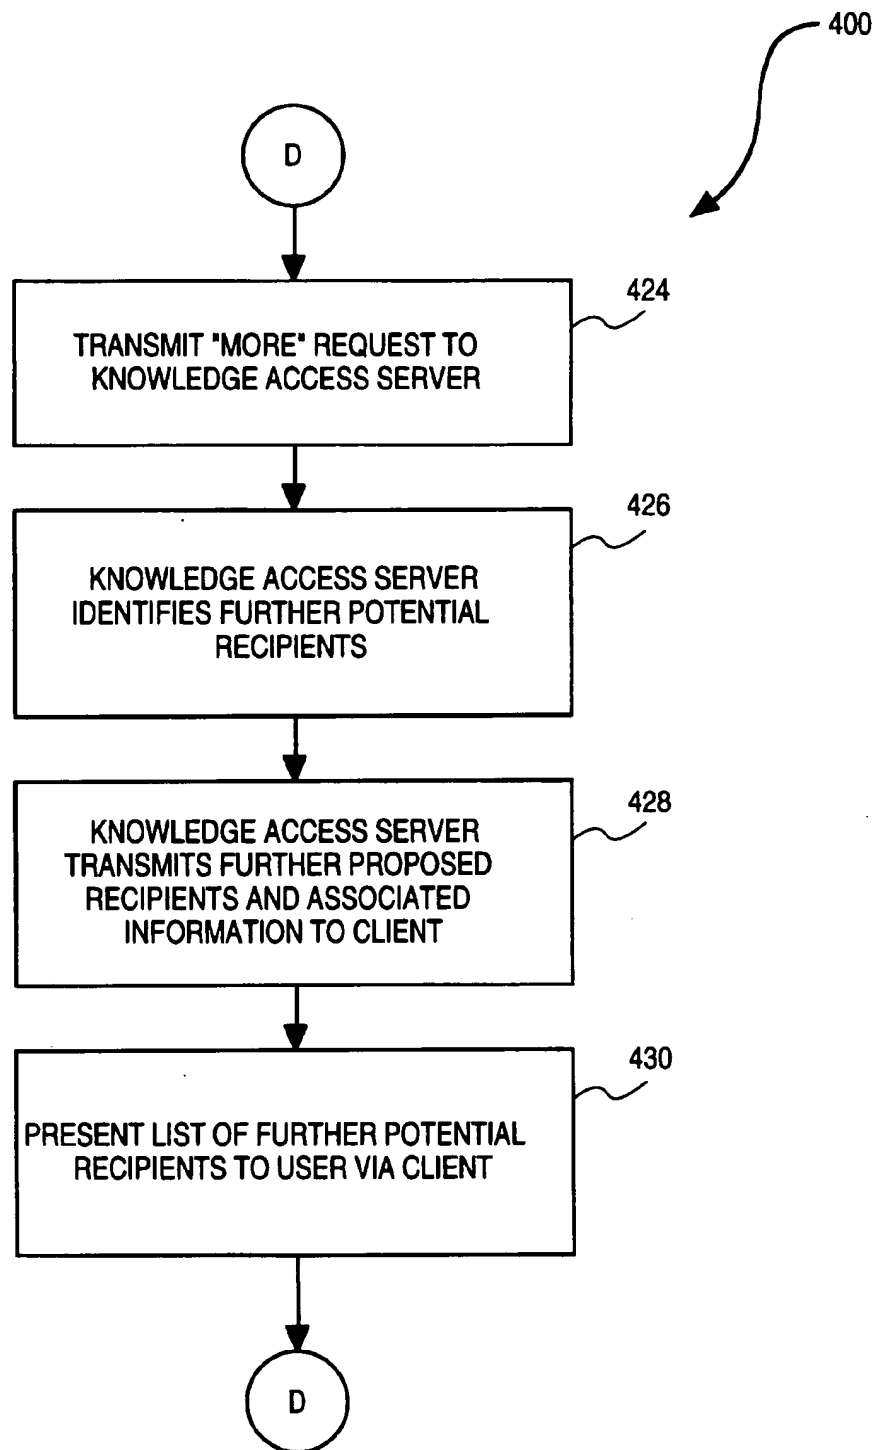


FIG. 17E



**FIG. 18B**

**FIG. 18C**




FIG. 18D

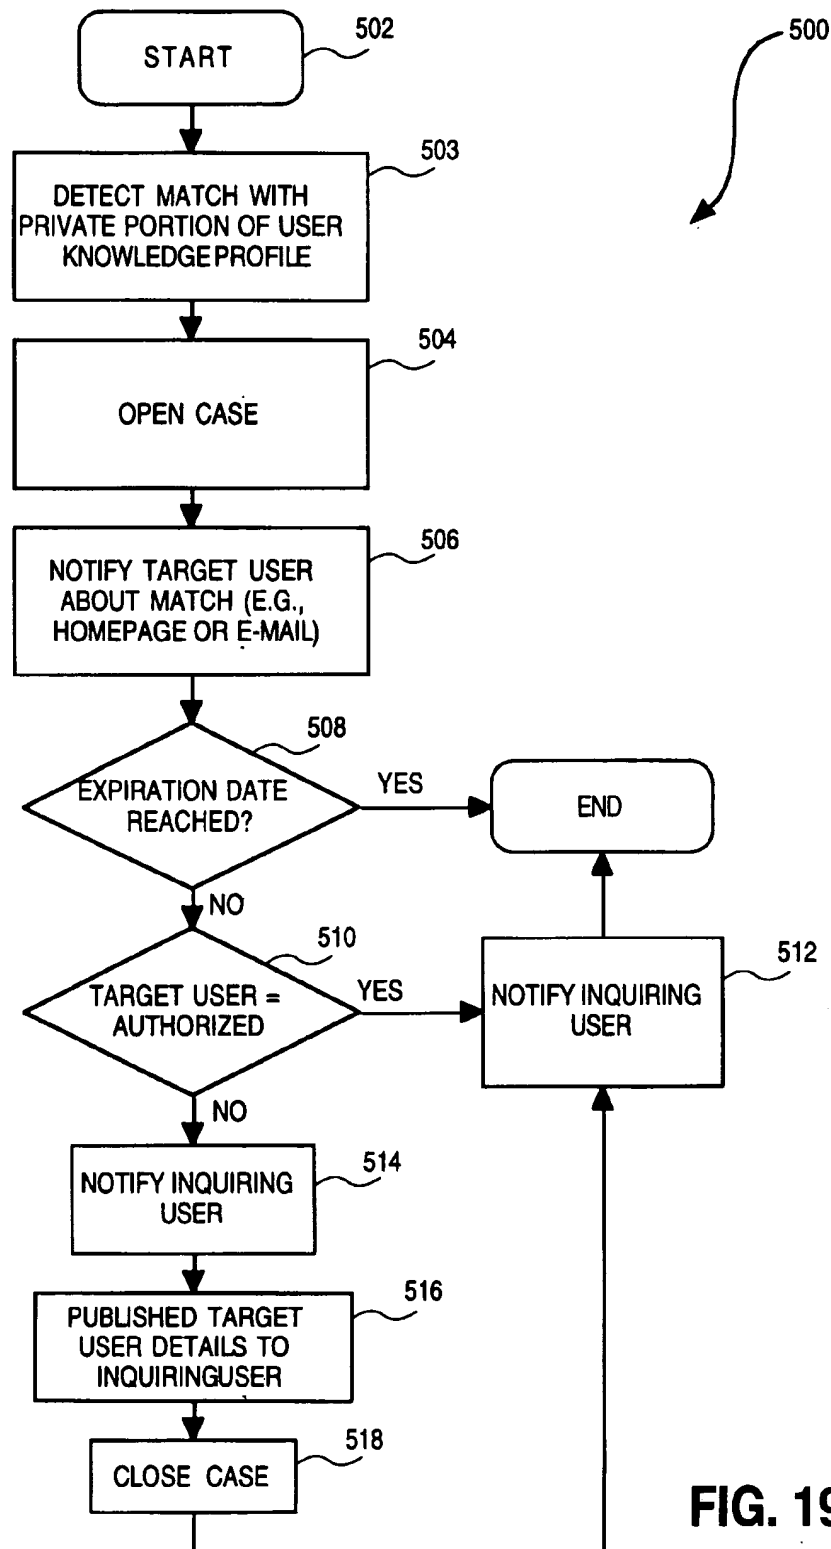


FIG. 19

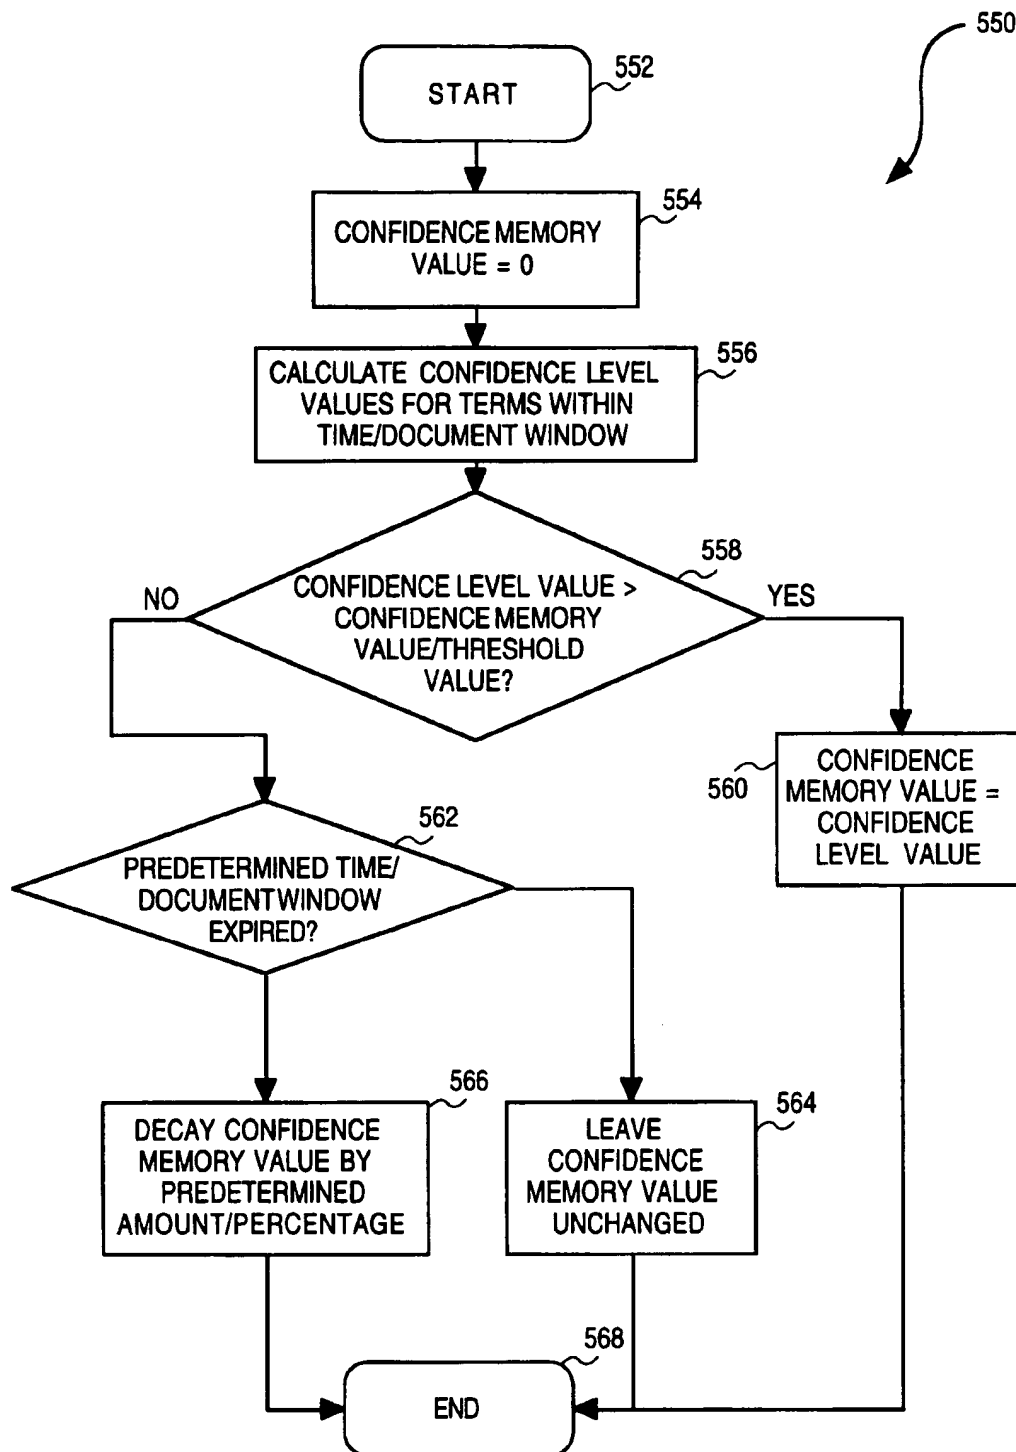


FIG. 20

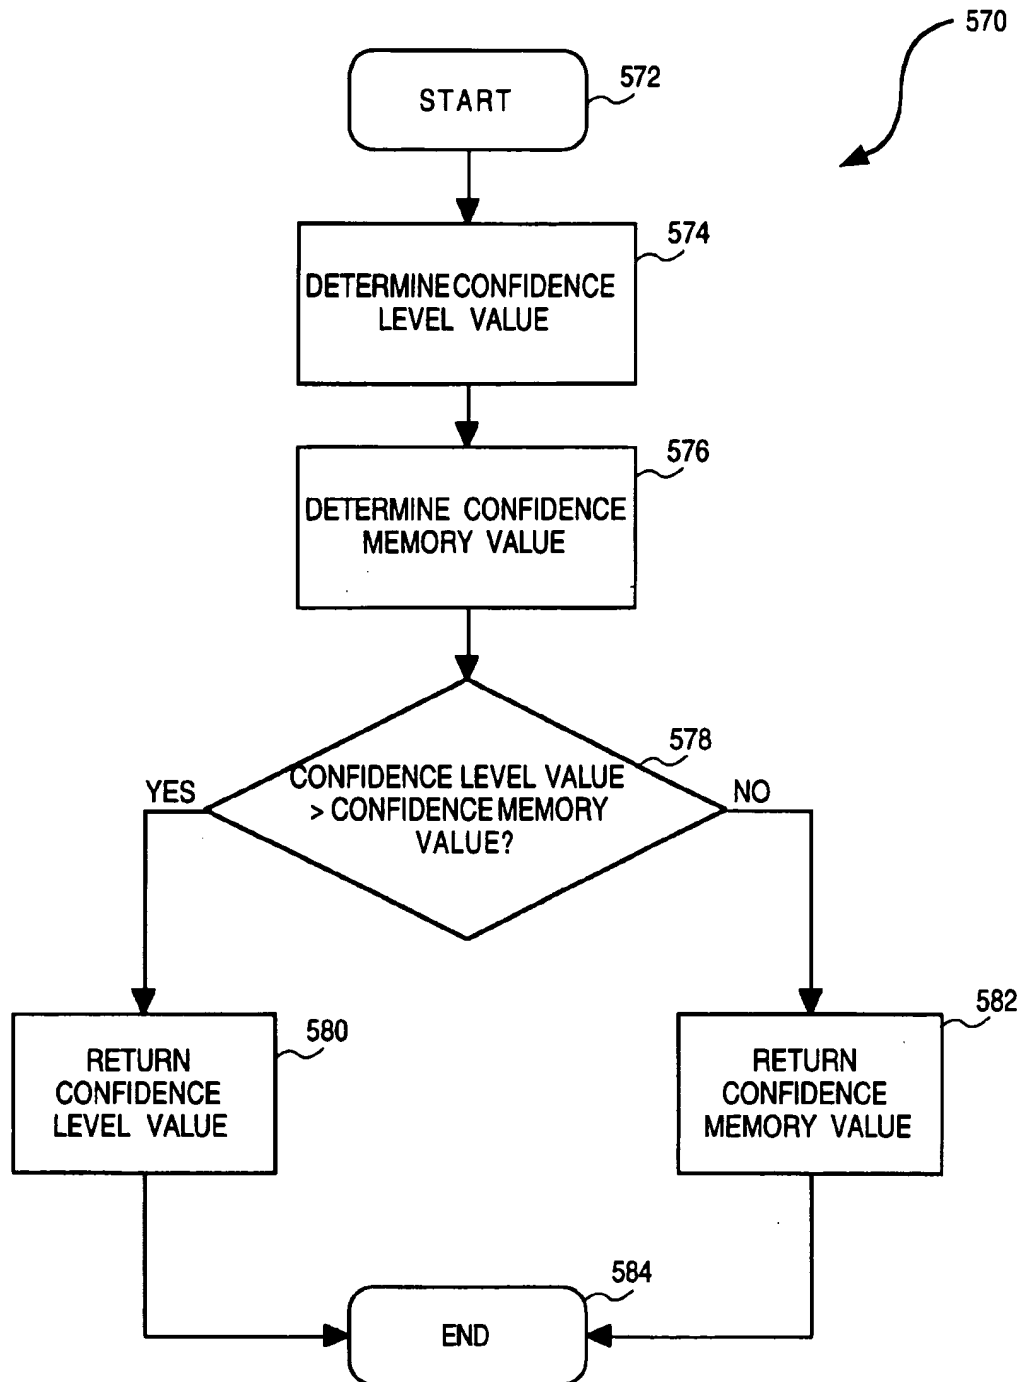


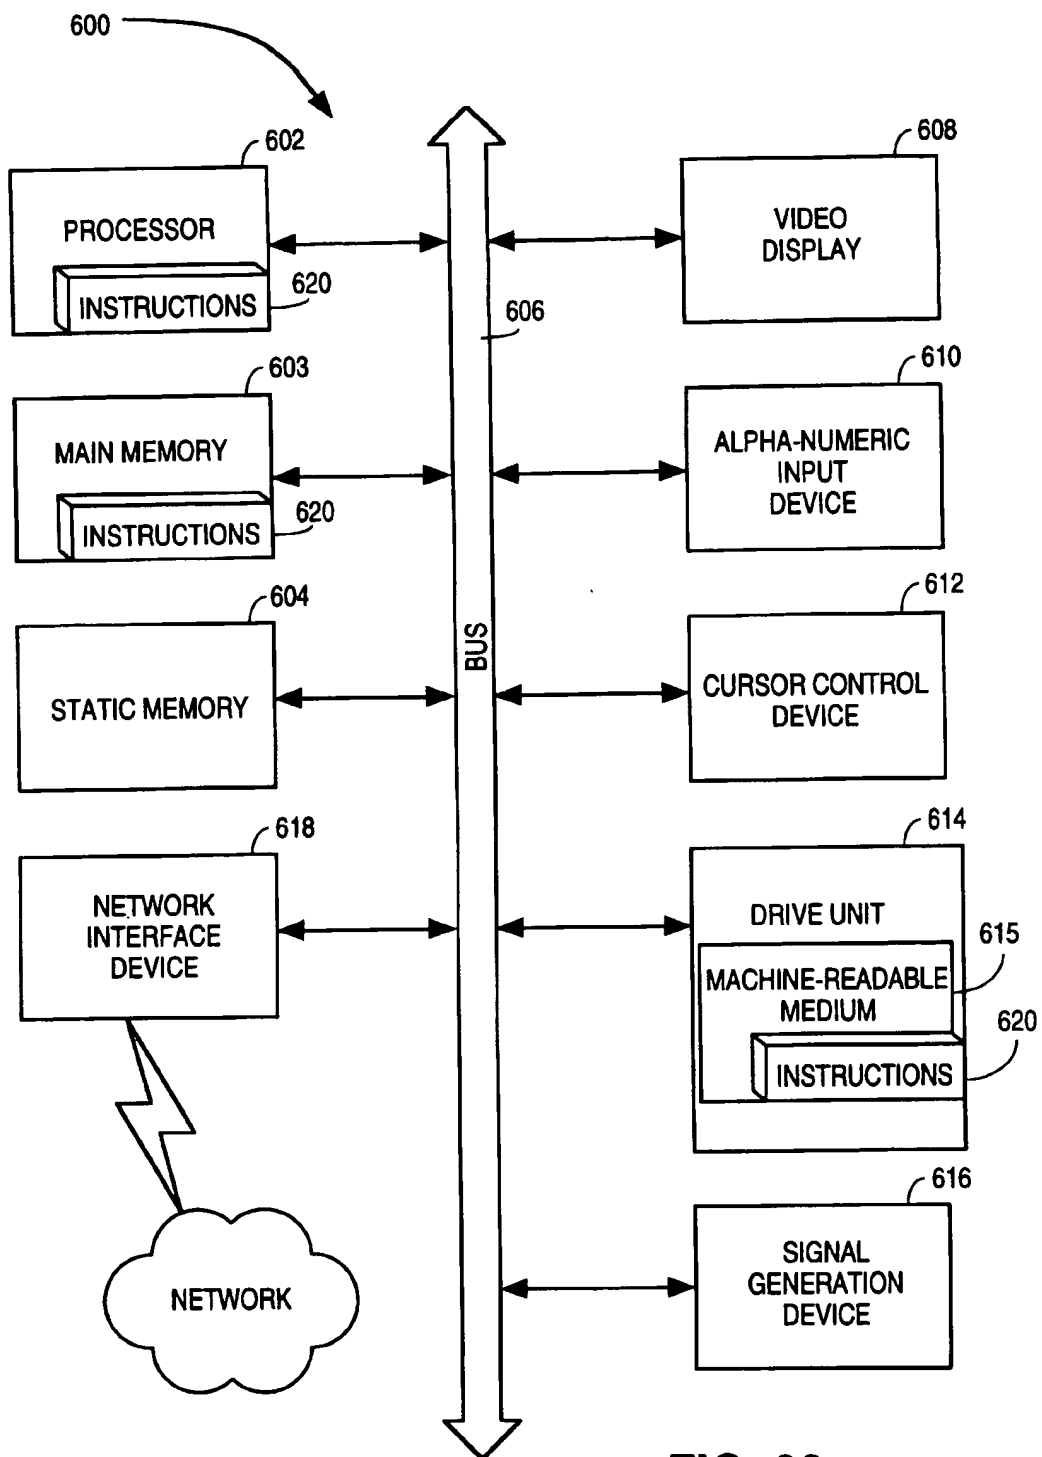
FIG. 21

112

USER ID	TERM ID	CONFIDENCE LEVEL	CONFIDENCE MEMORY	TIME STAMP
1	101	990	990	9812120800
1	106	976	976	9807241400
1	207	950	950	9805121200
.
.
.
1	67	310	310	9812241900
1	98	300	300	9802110700
1	116	290	290	9804120900
1	190	287	287	9806060900
.
.
.
8	107	995	995	9811191000
8	128	992	992	9811201000
8	154	900	900	9712291300
.
.
.
8	62	370	370	9812181421
8	8	300	300	9812211600
8	52	207	207	9802021647
8	261	200	200	9802221147
.
.
.

92 94 118 121 123

FIG. 22

**FIG. 23**

1

**METHOD AND SYSTEM FOR
CONSTRUCTING A KNOWLEDGE PROFILE
OF A USER HAVING UNRESTRICTED AND
RESTRICTED ACCESS PORTIONS
ACCORDING TO RESPECTIVE LEVELS OF
CONFIDENCE OF CONTENT OF THE
PORTIONS**

FIELD OF THE INVENTION

The present invention relates generally to the field of knowledge management and, more specifically, to a method and apparatus for automatically constructing a user knowledge profile and knowledge repository of electronic documents.

BACKGROUND OF THE INVENTION

The new field of "knowledge management" (KM) is receiving increasing recognition as the gains to be realized from the systematic effort to store and export vast knowledge resources held by employees of an organization are being recognized. The sharing of knowledge broadly within an organization offers numerous potential benefits to an organization through the awareness and reuse of existing knowledge, and the avoidance of duplicate efforts.

In order to maximize the exploitation of knowledge resources within an organization, a knowledge management system may be presented with two primary challenges, namely (1) the identification of knowledge resources within the organization and (2) the distribution and accessing of information regarding such knowledge resources within the organization.

The identification, capture, organization and storage of knowledge resources is a particularly taxing problem. Prior art knowledge management systems have typically implemented knowledge repositories that require users manually to input information frequently into pre-defined fields, and in this way manually and in a prompted manner to reveal their personal knowledge base. However, this approach suffers from a number of drawbacks in that the manual entering of such information is time consuming and often incomplete, and therefore places a burden on users who then experience the inconvenience and cost of a corporate knowledge management initiative long before any direct benefit is experienced. Furthermore, users may not be motivated to describe their own knowledge and to contribute documents on an ongoing basis that would subsequently be re-used by others without their awareness or consent. The manual input of such information places a burden on users who then experience the inconvenience and cost of a corporate knowledge management initiative long before any direct benefit is experienced.

It has been the experience of many corporations that knowledge management systems, after some initial success, may fail because either compliance (i.e., the thoroughness and continuity with which each user contributes knowledge) or participation (i.e., the percentage of users actively contributing to the knowledge management system) falls to inadequate levels. Without high compliance and participation, it becomes a practical impossibility to maintain a sufficiently current and complete inventory of the knowledge of all users. Under these circumstances, the knowledge management effort may never offer an attractive relationship of benefits to costs for the organization as a whole, reach a critical mass, and the original benefit of knowledge management falls apart or is marginalized to a small group.

2

In order to address the problems associated with the manual input of knowledge information, more sophisticated prior art knowledge management initiatives may presume the existence of a centralized staff to work with users to capture knowledge bases. This may however increase the ongoing cost of knowledge management and requires a larger up-front investment before any visible payoff, thus deterring the initial funding of many an otherwise promising knowledge management initiatives. Even if an initial decision is made to proceed with such a sophisticated knowledge management initiative, the cash expenses associated with a large centralized knowledge capture staff may be liable to come under attack, given the difficulty of quantifying knowledge management benefits in dollar terms.

As alluded to above, even once a satisfactory knowledge management information base has been established, the practical utilization thereof to achieve maximum potential benefit may be challenging. Specifically, ensuring that the captured information is readily organized, available, and accessible as appropriate throughout the organization may be problematic.

SUMMARY OF THE INVENTION

According to a first aspect of the invention, there is provided a method of constructing a user knowledge profile including first and second portions having different access restrictions. A confidence level is automatically assigned to content within an electronic document associated with a user, the content being potentially indicative of a user knowledge base. The content is then stored in either the first or the second portion of the user knowledge profile according to the assigned confidence level.

According to a second aspect of the invention, there is provided apparatus for constructing a user knowledge profile including first and second portions having different access restrictions. The apparatus includes confidence logic to examine an electronic document, associated with a user, and to assign a confidence level to content within the electronic document, the content being potentially indicative of a user knowledge base. The apparatus further includes a profiler to store the content in either the first or second portion of the user knowledge profile according to the assigned confidence level.

Other features of the present invention will be apparent from the accompanying drawings and from the detailed description that follows.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

FIG. 1 is a block diagram illustrating a knowledge management system, according to an exemplary embodiment of the present invention.

FIG. 2 is a block diagram illustrating a knowledge site management server, according to an exemplary embodiment of the present invention.

FIG. 3 is a block diagram illustrating a knowledge access server, according to an exemplary embodiment of the present invention.

FIG. 4 is a block diagram illustrating a knowledge converter, according to an exemplary embodiment of the present invention.

FIG. 5 is a block diagram illustrating a client software program, and an e-mail message generated thereby, according to an exemplary embodiment of the present invention.

FIG. 6 is a block diagram illustrating the structure of a knowledge repository, according to an exemplary embodiment of the present invention, as constructed from the data contained in a repository database and a user database.

FIG. 7 is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of constructing a user knowledge profile.

FIG. 8 is a flowchart illustrating a high-level method, according to an exemplary embodiment of the present invention, by which terms may be extracted from an electronic document and by which confidence level values may be assigned to such terms.

FIG. 9A is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of determining a confidence level for a term extracted from an electronic document.

FIG. 9B is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, by which a document weight value may be assigned to a document based on addressee information associated with the document.

FIG. 10 illustrates a term-document binding table, according to an exemplary embodiment of the present invention.

FIG. 11 illustrates a weight table, according to an exemplary embodiment of the present invention.

FIG. 12 illustrates an occurrence factor table, according to an exemplary embodiment of the present invention.

FIG. 13 illustrates a confidence level table, including initial confidence level values, according to an exemplary embodiment of the present invention.

FIG. 14 illustrates a modified confidence level table, including modified confidence level values, according to an exemplary embodiment of the present invention.

FIG. 15A is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of constructing a user knowledge profile that includes first and second portions.

FIG. 15B is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of storing a term in either a first or a second portion of a user knowledge profile.

FIG. 16A illustrates a user-term table, constructed according to the exemplary method illustrated in FIG. 15A.

FIG. 16B illustrates a user-term table, constructed according to the exemplary method illustrated in FIG. 15A.

FIG. 17A is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of facilitating access to a user knowledge profile.

FIG. 17B is a flowchart illustrating an alternative method, according to an exemplary embodiment of the present invention, of facilitating access to a user knowledge profile.

FIG. 17C is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of performing a public profile process.

FIG. 17D is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of performing a private profile process.

FIG. 17E is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of performing a profile modification process.

FIG. 18A is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of addressing an electronic document for transmission over a computer network.

FIG. 18B is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of executing an "explain" function that provides the reasons for the proposal of an e-mail recipient.

FIG. 18C is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of executing a "more" function that proposes further potential recipients for an e-mail message.

FIG. 18D illustrates a user dialog, according to an exemplary embodiment of the present invention, through which a list of potential recipients is displayed to an addressor of an e-mail message.

FIG. 19 is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of managing user authorization to publish, or permit access to, a user knowledge profile.

FIG. 20 is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of assigning a confidence value, either in the form of a confidence level value or a confidence memory value, to a term.

FIG. 21 is a flowchart illustrating a method, according to an exemplary embodiment of the present invention, of determining or identifying a confidence value, either in the form of a confidence level value or a confidence memory value, for a term.

FIG. 22 illustrates a user-term table, according to an exemplary embodiment of the present invention, that is shown to include a confidence level value column, a confidence memory value column and a time stamp column.

FIG. 23 is a block diagram illustrating a machine, according to one exemplary embodiment, within which software in the form of a series of machine-readable instructions, for performing any one of the methods discussed above, may be executed.

DETAILED DESCRIPTION

A method and apparatus for constructing and maintaining a user knowledge profile are described. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be evident, however, to one skilled in the art that the present invention may be practiced without these specific details.

OVERVIEW

With a view to addressing the above described difficulties associated with manual knowledge capture either by a profile owner or by a dedicated staff, there is provided a method and apparatus for capturing knowledge automatically, without excessive invasion or disruption of normal work patterns of participating users. Further, the present specification teaches a method and apparatus whereby a database of captured knowledge information is maintained continuously and automatically, without requiring that captured knowledge information necessarily be visible or accessible to others. The present specification also teaches facilitating the user input and modification of a knowledge profile associated with the user in a knowledge database at the user's discretion.

The present specification teaches a method and apparatus for intercepting electronic documents, such as for example e-mail messages, originated by a user, and extracting terms therefrom that are potentially indicative of a knowledge base of the originating user. The extracted knowledge terms may then be utilized to construct a user knowledge profile. The

5

grammatical structure, length, frequency and density with which the extracted knowledge terms occur within electronic documents originated by a user, and prior history of use of the extracted knowledge terms within an organization may furthermore be utilized to attach a metric, in the form of a confidence level value, to the relevant knowledge terms for the purpose of grouping, ranking, and prioritizing such knowledge terms. Knowledge terms may furthermore be stored in either a private or public portion of the user knowledge profile, depending upon the confidence level values thereof.

It will be appreciated that the large volume of e-mail messages traversing an e-mail system over a period of time will contain a large number of terms that may be irrelevant to the identification of the knowledge base of a user. With a view to determining which terms are truly indicative of a knowledge base, a number of rules (or algorithms) may be exercised with respect to extracted terms to identify terms that are candidates for inclusion within a public portion of the user knowledge profile. Further rules (or algorithms) may be applied to an assembled knowledge profile for the purpose of continually organizing and refining the profile.

Corporate e-mail systems have become increasingly pervasive, and have become an accepted medium for idea communication within corporations. Accordingly, the content of e-mail messages flowing within a large organization amounts to a vast information resources that, over the course of time, may directly or indirectly identify knowledge bases held by individuals within the organization.

The present specification also teaches addressing privacy concerns associated with the examination of e-mail messages for the above purposes by providing users with the option selectively to submit originated e-mail messages for examination, or alternatively to bypass the examination and extraction system of the present invention.

There is also taught a computer-implemented method and apparatus for addressing an electronic document, such as an e-mail message, for transmission over a computer network. The e-mail message may be examined to identify terms therein. The identified terms are then compared to a number of user knowledge profiles with a view to detecting a predetermined degree of correspondence between the identified terms and any one or more of the user knowledge profiles. In the event that a predetermined degree of correspondence is detected, the sender of the electronic document is prompted to the either accept or decline the proposed recipient as an actual recipient of the electronic document, after first being offered an opportunity to inspect the specific basis of the correspondence between the identified terms and the proposed recipients. The e-mail message may also be parsed to extract recipients entered manually by the user. The degree of correspondence between the knowledge profiles of the manually entered recipients and the identified terms of the message is then optionally used as the basis of recommendations to the user that certain manually entered recipients be dropped from the ultimate list of recipients.

This aspect of the present teachings is advantageous in that a sender of an e-mail message is presented with a list of proposed recipients, identified according to their knowledge profiles and the content of the e-mail message, who may be interested in receiving the e-mail message. Accordingly, the problems of over-distribution and under-distribution of e-mail messages that may be encountered within an organization may be reduced. Specifically, in the over-distribution situation, many users are frequently copied on e-mail messages, resulting in lost productivity as the users

6

struggle to cope with increasing volumes of daily e-mail. Further, when the time available to read e-mail messages becomes restricted, users typically begin to defer reading of e-mail messages, and communication efficiency within the organization may be adversely affected. In the under-distribution situation, it may occur that the proper recipients of the message are not included in the distribution list, and accordingly fall "out of the loop".

There is also taught a method of facilitating a user profile query or look-up wherein, in response to a match between a query and a user profile, the owner of the user profile may be prompted for authorization to publish all (or a portion) of the user profile to the originator of the query or to others generally. This is advantageous in that it addresses the above mentioned privacy concerns by treating the knowledge profile as a confidential resource under the control of the user. The user is thus also able to control the timing, circumstances and extent to which it is made accessible to others. A further advantage is that the user is prompted for input specifically to satisfy specific, pending requests of others. This relieves the user of the need to remember to modify his or her profile on a regular basis and the need to make decisions concerning the composition of the profile prospectively, prior to any actual use of the profile by others. In this manner the user saves time and effort, since the determination that manual interaction with the profile is necessary is a function of the present system, not a responsibility of the user.

There is also taught a method of assigning a confidence level value to a term within an electronic document. This confidence level value is based on a first quantitative indicator, derived from the number of occurrences of the term within the electronic document, and a second characteristic indicator, derived utilizing the characteristic of the term.

For the purposes of the present application, the word "term" shall be taken to include any acronym, word, collection of words, phrase, sentence, or paragraph. The term "confidence level" shall be taken to mean any indication, numeric or otherwise, of a level within a predetermined range.

SYSTEM ARCHITECTURE

FIG. 1 is a block diagram illustrating a knowledge management system 10, according to an exemplary embodiment of the present invention. The system 10 may conveniently be viewed as comprising a client system 12 and a server system 14. The client system 12 may comprise one or more clients, such as browser clients 16 and e-mail clients 18, that are resident on terminals or computers coupled to a computer network. In one exemplary embodiment, each of the browser clients 16 may comprise the Internet Explorer client developed by Microsoft Corp. of Redmond, Wash., or the Netscape Navigator client developed by Netscape Communications of Menlo Park, Calif. Each of the e-mail clients 18 may further comprise the Outlook Express, Outlook 97, Outlook 98 or Netscape Communicator e-mail programs. As will be described in further detail below, the browser and e-mail clients 16 are complemented by extensions 19, that enable the e-mail clients 18 to send an electronic message (e.g., either an e-mail or HTML document) to a knowledge server 22 implemented on the server side 14 of the system 10. As shown in FIG. 1, the extensions 19 may be integral with an e-mail client 18, or external to the client 18 and in communication therewith. The clients 16 and 18 may default to sending every communication to a relevant component of

the knowledge server 22, while allowing a user specifically to designate a communication not suitable for transmission to the knowledge server 22. The user designation may be facilitated through controls that are installed as software modules which interact with or modify an e-mail client 18, and which cause messages to be copied to a special e-mail address (e.g., a Knowledge Server (KS) mailbox 25 maintained by a e-mail server 23) associated with a knowledge server component. In the case where a client extension 19 for performing this automatic transmission is not available, the user can manually add the e-mail address of the KS mailbox 25 to the list of recipients for the message. Further details in this regard are provided below. Files embedded within an e-mail message, such as attachments, may also be selectively included or excluded from the capture process and may also be selectively included or excluded from retention in a knowledge repository.

The browser clients 16 are used as an additional means to submit documents to the knowledge server 22 at the discretion of a user. The browser client 16 is used to access an interface application 34, maintained on a web server 20, which transmits documents to the knowledge server 22.

In alternate embodiments, a client may also propagate a list of bookmarks, folders or directories to the knowledge server 22 for the purpose of user knowledge profile construction.

SERVER SIDE ARCHITECTURE

The server side 14 of the system 10 includes the web server 20, the e-mail server 23 and the knowledge server 22. The web server 20 may be any commercially available web server program such as Internet Information Server (IIS) from Microsoft Corporation, the Netscape Enterprise Server, or the Apache Server for UNIX. The web server 20 includes the interface application 34 for interfacing with the knowledge server 22. The web server 20 may run on a single machine that also hosts the knowledge server 22, or may alternatively run along with the interface application 34 on a dedicated web server computer. The web server 20 may also be a group of web server programs running on a group of computers to thus enhance the scalability of the system 10. As the web server 20 facilitates access to a local view of a knowledge repository 50, maintained by the knowledge access server 26, by the browser clients 16, the web server interface application 34 implements knowledge application interfaces, knowledge management interfaces, user profile creation and maintenance interfaces, and a server management interface. The web server 20 also facilitates knowledge profile queries, e-mail addressing to an e-mail client 18, and any other access to the knowledge server 22 using the standard HTTP (web) protocol.

The knowledge server 22 includes a knowledge site management server (KSMS) 27 and the knowledge access server (KAS) 26. The knowledge server access 26 includes an interface that provides a local view of a knowledge repository 50, which is physically stored in the user database 56A and a repository database 56B. The knowledge site management server 27 is shown to have access to the local view of the knowledge repository 50 maintained by the knowledge access server 26. The illustrated components of the knowledge server 22 are collectively responsible for the capture (termed "knowledge discovery") of terms indicative of a user knowledge base and for the distribution of user knowledge profile information. Knowledge discovery may be done by the examination and processing of electronic documents, such as e-mail messages, which may be propa-

gated to the e-mail server 23 from an e-mail client 18 via the Simple Mail Transfer Protocol (SMTP), as shown at 32. Alternatively, knowledge discovery may be implemented by the examination of submissions from a browser client 16 via the web server 20.

The knowledge server 22 includes the knowledge access server 26 and the knowledge site management server 27 as two separate and distinct server systems in view of the divergent functions provided by the servers 26 and 27. Specifically, the knowledge site management server 27 functions primarily to manage non-interactive processing (e.g., the extraction of knowledge from inbound e-mail messages), to manage the user information database 56A, and to implement various centralized system management processes. The knowledge site management server 27 does not communicate interactively with clients 18, or with clients 16 except for administrative functions. The knowledge access server 26, on the other hand, functions primarily to respond to queries and updates from users submitted via clients, typically browser clients 16. Multiple instances of a knowledge access server 26 may be required to support a large corporate environment and to provide appropriate scalability; however only one knowledge site management server 27, one user database 56A, and one repository database 56B typically exist in a working system. In small scale environments, the web server 20, knowledge access server 26, and knowledge site management server 27, and even the e-mail server 23, may all optionally be deployed on the same physical computer.

FIG. 2 is a block diagram illustrating an exemplary embodiment, according to the present invention, of the knowledge site management server 27. The server 27 is shown to include a socket front-end 40 to facilitate communication with the web server 20 for administrative requests, a request handler 44, a knowledge gathering system 28, a knowledge converter 24, and a variety of specialized controller modules 45A-45C. The request handler 44, upon receiving a request from the web server 20 via the interface application 34 and socket front-end 40, starts a session to process the request such as, for example, a request by an authorized systems administrator to configure the behavior of the knowledge gathering system 28.

The knowledge gathering system 28 is shown in FIG. 2 to include an extraction controller 47, a mail system interface 42, and a term extractor 46 including confidence logic 45. The extraction controller 47 commands the mail system interface 42 to retrieve messages submitted by the e-mail client extensions 19 to the KS mailbox 25 on the e-mail server 23 for the purpose of extraction and processing. The extraction controller 47 can request this continuously or periodically on a scheduled basis, so that messages can be processed at a convenient time when computing resources are lightly loaded, for example, overnight. The mail system interface 42 retrieves e-mail messages from the e-mail server 23 using the Simple Mail Transfer Protocol (SMTP), Post Office Protocol 3 (POP3), or Internet Message Access Protocol 4 (IMAP4) protocols. The mail system interface 42 propagates electronic documents directly to a term extractor 46, including confidence logic 45, that operates to convert electronic documents into per-user knowledge profiles that are stored in a knowledge repository 50. The term extractor 46 may include any commercially available term extraction engine (such as "NPTOOL" from LingSoft Inc. of Helsinki, Finland, or "Themes" from Software Scientific) that analyzes the electronic document, recognizes noun phrases in the document, and converts such phrases to a canonical form for subsequent use by the confidence logic 45 as candidate terms in a knowledge profile.

The term extractor 46 performs a variety of the steps when parsing and decoding an electronic document, such as interpreting any special attributes or settings encoded into the header of the message of the e-mail client 18, resolving the e-mail addresses of recipients against either the built-in user database or an external user database, preprocessing the electronic document, extracting noun-phrases from the text as candidates for knowledge terms, processing these knowledge terms, and storing summary information about the document and extraction process in the databases 56A and 56B. The term extractor 46 further detects and strips out non-original texts, attachments and in some cases the entire electronic document based on the document not meeting predetermined minimum criteria. Further details regarding the exact procedures implemented by the term extractor 46 will be provided below. Once the term extractor 46 has extracted the knowledge terms, the knowledge repository 50 is updated. Specifically, new terms are added, and repetitions of known terms are used to update the knowledge repository 50.

1b) The knowledge repository 50 is defined by a hierarchical structure of classes. The objects of these classes represent the knowledge information that includes, inter alia, user profiles (including knowledge profiles) and organizational structure, and are stored in two databases: the user database 56A and the repository database 56B. The repository database 56B contains profile and repository information and can use one of a number of commercial relational database management systems that support the Open DataBase Connectivity (ODBC) interface standard. A database interface 54 provides a logical database-independent class API to access the physical databases and to shield the complete server codes from accessing database native API so that the server process can use any relational database management system (RDMS). Because the repository database 56A is open to inspection by systems administrators, and may be hosted on an existing corporate system, special measures may be taken to enhance the privacy of information in the repository database 56B; for example, the repository database 56B contains no actual user names or e-mail addresses, but instead may use encrypted codes to represent users in a manner that is meaningful only in combination with the user database. The user database 56A is a small commercial RDBMS embedded into the knowledge repository 50 in such a way that it cannot be accessed except through the interfaces offered by the system 10. The user database 56A contains encrypted identifying codes that allow the names of actual users to be associated with e-mail addresses, login IDs, passwords, and profile and repository information in the repository database.

A lexicon controller 45C is responsible for building tables of associated terms. Terms are considered "associated" with each other to the extent that they tend to co-occur in close proximity within the documents of multiple users. The lexicon controller 45C manages the background process of data mining that is used to discover associations between terms and record those in special association tables within the repository database 56B.

A profile controller 45B is a module that may optionally be included within the knowledge site management server 27, and manages a queue of pending, compute-intensive operations associated with updating profiles. Since the algorithm for the confidence level value calculation of a term (embodied in the confidence logic 45) depends on the total number of documents profiled, the confidence level value for each and every term in a user's profile is technically obsolete when any document is profiled. The profile con-

troller 45B manages the "recalculation" of profiles. The actual operation is performed within the knowledge access server 26, which has a knowledge repository 50 interface.

1a) A case controller 45A keeps track of open cases and initiates notifications to users concerning their status. A "case" is a pending request from one user to another, as will be detailed below. For example, if a user requests an expert in a certain field via a client browser client 16, the knowledge access server 26 matches the term against both the public and private portions of all user profiles. If a high confidence, but private, match is found, the system cannot reveal the identity of the matched person to the inquirer and must therefore open a "case". The case places a notification in the profile "home" page of the target user and/or transmits an e-mail message with a link back to that page. The target user may then (via a browser):

1. See the identity of the inquirer and the basis of the match.
2. See comments added by the inquirer.
3. Deny the request, at which point the case is closed.
4. Put a block on any further matches from that person or based on that term.
5. Go into the profile and edit the term responsible for the match.
6. Indicate that the case is accepted and provide authorization to reveal the identity of the target to the inquirer.

2c) From the perspective of the inquirer, private matches are initially returned with a match strength only and do not reveal the name of the person or document matched. The user can then initiate cases for any or all of these private matches, based on how urgently the information is needed, how good the matches were, and whether the public matches are sufficient. Each case gets an expiration date set by the inquirer and notification options regarding how the inquirer wants to be told about the disposition of the case. Open cases are summarized in the Web area for the inquirer, along with the date and query that generated the return values. If the target denies a case, that status is communicated to the user. The user has no option to send e-mail or otherwise further identify that person. If the target accepts the case, the identity of the target is communicated to the user by updating the case record and the case is closed. Case history retention options are a site administration option.

FIG. 3 is a block diagram illustrating the components that constitute the knowledge access server 26. The knowledge access server 26 is shown to include a socket front-end 40 to facilitate communication with the web server interface application 34. The knowledge access server 26 further includes a request handler 44, a term extractor 46, a knowledge repository 50 and a database interface 54 that function in a manner similar to that described above with reference to the knowledge gathering system 28. The term extractor 46 includes comparison logic 51, the functioning of which will be described below. The knowledge access server 26 functions primarily as an interface between knowledge users and the knowledge repository 50. It provides services to the web server interface application 34, which implements a number of user interfaces as described above for interacting with the knowledge repository 50.

FIG. 4 is a block diagram illustrating the components that constitute the knowledge converter 24. The knowledge converter 24 is shown to include a term extractor 46 that is fed from an array of format converters 60. The knowledge converter 24 is able to access the knowledge repository 50, and to import data from other knowledge systems, or export knowledge to other knowledge systems, via each of the format converters 60.

Returning to FIG. 1, the knowledge access server 26 implements the interface to the knowledge repository 50 and

11

the knowledge site management server 27 is shown to access the knowledge repository 50 via the knowledge access server 26. FIGS. 3 and 4 illustrate data for the knowledge repository 50 as residing in databases 56A and 56B. The databases 56A and 56B are built on a general database interface 54 and provide persistent storage for the core system classes referred to above. In one exemplary embodiment of the present invention, the user database and the repository databases are implemented utilizing the Microsoft SQL server, developed by Microsoft Corp. of Redmond Wash., to provide default storage management services for the system. However, programming may be done at a more general level to allow for substitution of other production class relational database management systems, such as those developed by Sybase, Oracle or Informix.

CLIENT SIDE ARCHITECTURE

FIG. 5 is a diagrammatic representation of a client, according to an exemplary embodiment of the present invention, in the form of an e-mail client 18. It will be appreciated that the e-mail client 18 may be any commercially available e-mail client, such as a Microsoft Exchange, Outlook Express, Outlook 97/98 or Lotus Notes client. The e-mail client 18 includes modifications or additions, in the form of the extensions 19, to the standard e-mail client to provide additional functionality. Specifically, according to an exemplary embodiment of the present invention, three subsystems are included within the e-mail client extensions 19, namely a user interface 80, a profiling system 82, and an addressing system 84.

The profiling system 82 implements properties on an originated message, as well as menu and property sheet extensions at global and message levels for users to set and manipulate these new properties. More specifically, profiling system 82 provides a user with a number of additional options that determine how a message 85 propagated from the e-mail client 18 to the knowledge repository 50 will be processed and handled for the purposes of knowledge management. A number of the provided options are global, while others apply on a per-message basis. For example, according to one exemplary embodiment, the following per-message options (or flags) may be set by a user to define the properties of an e-mail message:

1. An "Ignore" flag 86 indicating the e-mail message should not be processed for these purposes of constructing or maintaining a user knowledge profile, and should not be stored.
2. A "Repository" parameter 88 indicating that the message may be processed for the purposes of constructing a knowledge profile and then stored in the repository 50 for subsequent access as a document by others. The "Repository" parameter 88 also indicates whether the document (as opposed to terms therein) is to be stored in a private or public portion of the repository 50.

A number of global message options may also be made available to a user for selection. For example, an e-mail address (i.e., the KS mailbox 25 or the e-mail server 23) for the knowledge server 22 may be enabled, so that the e-mail message is propagated to the server 22.

Actual implementation and presentation of the above per-message and global options to the user may be done by the addition of a companion application or set of software modules which interact with API's provided by e-mail clients, or modules which modify the e-mail client itself, which are available during message composition. If the user activates the Ignore flag 86, the profiling system 82 will not

12

make any modifications to the message and no copy of the message will be sent to the knowledge gathering system 28 via the KS mailbox 25. Otherwise, per-message options, once obtained from the user, are encoded. Subsequently, when the user chooses to send the message 85 using the appropriate control on the particular e-mail client 18, the e-mail address of the knowledge gathering server is appended to the blind copy list for the message. The profiling system 82 encrypts and encodes the following information into the message header, for transmission to and decoding by the knowledge gathering system 28, in accordance with Internet specification RFC 1522:

1. The list of e-mail addresses in the "to:" and "cc:" lists;
2. Per-message options as appropriate; and
3. For those recipients suggested by the addressing system 84 (see below), a short list of topic identifiers including the primary topics found within the message and the primary topics found within the user profile that formed a basis of a match.

4. Security information to validate the message as authentic.

When the message 85 is sent over the normal e-mail transport, the following events occur:

1. Recipients on the "to:" and "cc:" lists will receive a normal message with an extra header containing the encoded and encrypted options. This header is normally not displayed by systems that read e-mail and can be ignored by recipients;
2. The recipients will not be aware that the knowledge gathering system has received a blind copy of the message; and
3. If the sender chooses to archive a copy of the message 85, the e-mail address of the knowledge gathering system 28 will be retained in the "bcc" field as a reminder that the message was sent to the knowledge gathering server.

Further details concerning the addressing system 86 will be discussed below.

THE REPOSITORY

FIG. 6 is a block diagram illustrating the structure of the repository 50, according to one exemplary embodiment of the present invention, as constructed from data contained in the repository database 56B, and the user database 56A. The repository 50 is shown to include a number of tables, as constructed by a relational database management system (RDBMS). Specifically, the repository 50 includes a user table 90, a term table 100, a document table 106, a user-term table 112, a term-document table 120 and a user-document table 130. The user table 90 stores information regarding users for whom knowledge profiles may be constructed, and includes an identifier column 92, including unique keys for each entry or record within the table 90. A name column 94 includes respective names for users for whom knowledge profiles are maintained within the repository 50. A department column 96 contains a description of departments within an organization to which each of the users may be assigned, and an e-mail column 98 stores respective e-mail addresses for the users. It will be appreciated that the illustrated columns are merely exemplary, and a number of other columns, storing further information regarding users, may be included within the user table 90.

The term table 100 maintains a respective record for each term that is identified by the term extractor 46 within an electronic document, and that is included within the repository

tory 50. The term table 100 is shown to include an identifier column 102, that stores a unique key for each term record, and a term column 104 within which the actual extracted and identified terms are stored. Again, a number of further columns may optionally be included within the term table 100. The document table 106 maintains a respective record for each document that is processed by the term extractor 46 for the purposes of extracting terms therefrom. The document table 106 is shown to include an identifier column 108, that stores a unique key for each document record, and a document name column 110, that stores an appropriate name for each document analyzed by the term extractor 46.

The user-term table 112 links terms to users, and includes at least two columns, namely a user identifier column 114, storing keys identifying users, and a term identifier column 116, storing keys identifying terms. The user-term table 112 provides a many-to-many mapping of users to terms. For example, multiple users may be associated with a single term, and a single user may similarly be associated with multiple terms. The table 112 further includes a confidence level column 118, which stores respective confidence level values, calculated in the manner described below, for each user-term pair. The confidence level value for each user-term pair provides an indication of how strongly the relevant term is coupled to the user, and how pertinent the term is in describing, for example, the knowledge base of the relevant user.

The term-document table 120 links terms to documents, and provides a record of which terms occurred within which document. Specifically, the term-document table 120 includes a term identifier column 122, storing keys for terms, and a document identifier column 124, storing keys for documents. The table 120 further includes an adjusted count column 126, which stores values indicative of the number of occurrences of a term within a document, adjusted in the manner described below. For example, the first record within the table 120 records that the term "network" occurred within the document "e-mail 1" 2.8 times, according to the adjusted count.

The user-document table 130 links documents to users, and includes at least two columns, namely a user identifier column 132, storing keys identifying users, and a document identifier column 134, storing keys identifying various documents. For example, the first record within the exemplary user-document table 130 indicates that the user "Joe" is associated with the document "e-mail 1". This association may be based upon the user being the author or recipient of the relevant document.

IDENTIFICATION OF KNOWLEDGE TERMS AND THE CALCULATION OF ASSOCIATED CONFIDENCE LEVEL VALUES

FIG. 7 is a flow chart illustrating a method 140, according to an exemplary embodiment of the present invention, of constructing a user knowledge profile. FIG. 7 illustrates broad steps that are described in further detail with reference to subsequent flow charts and drawings. The method 140 commences at step 142, and proceeds to decision box 144, wherein a determination is made as to whether an electronic document, for example in the form of an e-mail propagated from an e-mail client 18, is indicated as being a private document. This determination may be made at the e-mail client 18 itself, at the e-mail server 23, or even within the knowledge site management server 27. This determination may furthermore be made by ascertaining whether the Ignore flag 86, incorporated within an e-mail message 85, is

set to indicate the e-mail message 85 as private. As discussed above, the Ignore flag 86 may be set at a users discretion utilizing the profiling system 82, accessed via the user interface 80 within the extensions 19 to the e-mail client 18. In the event that the electronic document is determined to be private, the method 140 terminates at step 146, and no further processing of the electronic document occurs. Alternatively, the method 140 proceeds to step 148, where confidence level values are assigned to various terms within the electronic document. At step 150, a user knowledge profile is constructed utilizing the terms within the electronic document to which confidence level values were assigned at step 148. The method 140 then terminates at step 146.

FIG. 8 is a flow chart illustrating a high-level method 148, according to an exemplary embodiment of the present invention, by which terms may be extracted from an electronic document, and by which confidence level values may be assigned such terms. The method 148 comprises two primary operations, namely a term extraction operation indicated at 152, and a confidence level value assigning operation, indicated at a 154. The method 148 implements one methodology by which the step 148 shown in FIG. 7 may be accomplished. The method 148 begins at step 160, and then proceeds to step 162, where an electronic document, such as for example an e-mail, a database query, a HTML document and or a database query, is received at the knowledge site management server 27 via the mail system interface 42. For the purposes of explanation, the present example will assume that an e-mail message, addressed to the KS mailbox 25, is received at the knowledge site management server 27 via the mail system interface 42, from the e-mail server 23. At step 164, terms and associated information are extracted from the electronic document. Specifically, the e-mail message is propagated from the mail system interface 42 to the term extractor 46, which then extracts terms in the form of, for example, grammar terms, noun phrases, word collections or single words from the e-mail message. The term extractor 46 may further parse a header portion of the e-mail to extract information therefrom that is required for the maintenance of both the repository and user databases 56B and 56A. For example, the term extractor 46 will identify the date of transmission of the e-mail, and all addressees. The term extractor 46 will additionally determine further information regarding the electronic document and terms therein. For example, the term extractor 46 will determine the total number of words comprising the electronic document, the density of recurring words within the document, the length of each term (i.e., the number of words that constitute the term), the part of speech that each word within the document constitutes, and a word type (e.g., whether the word is a lexicon term). To this end, the term extractor 46 is shown in FIG. 2 to have access to a database 49 of lexicon terms, which may identify both universal lexicon terms and environment lexicon terms specific to an environment within which the knowledge site management server 27 is being employed. For example, within a manufacturing environment, the collection of environment lexicon terms will clearly differ from the lexicon terms within an accounting environment.

Following the actual term extraction, a first relevancy indicator in the form of an adjusted count value is calculated for each term within the context of the electronic document at step 168. At step 170, a second relevancy indicator in the form of a confidence level is calculated for each term within the context of multiple electronic documents associated with a particular user. Further details regarding steps 168 and 170 are provided below. The method 148 then terminates at step 172.

FIG. 9A is a flow chart illustrating a method 154, according to an exemplary embodiment of the present invention, of determining a confidence level for a term extracted from an electronic document. Following the commencement step 180, a term and associated information is received at the confidence logic 45, included within the term extractor 46. While the confidence logic 45 is shown to be embodied in the term extractor 46 in FIG. 2, it will be appreciated that the confidence logic 45 may exist independently and separately of the term extractor 46. In one embodiment, the associated information includes the following parameters:

1. A count value indicating the number of occurrences of the term within a single electronic document under consideration;
2. A density value, expressed as a percentage, indicating the number of occurrences of the term relative to the total number of terms within the electronic document;
3. A length of value indicating the total number of words included within the relevant term;
4. A Part of Speech indication indicating the parts of speech that words included within the term comprise (e.g., nouns, verbs, adjectives, or adverbs); and
5. A Type indication indicating whether the term comprises a universal lexicon term, an environment lexicon term, or is of unknown grammatical structure.

At step 184, a "binding strength", indicative of how closely the term is coupled to the electronic document under consideration, is determined. While this determination may be made in any number of ways, FIG. 10 shows an exemplary term-document binding table 200, utilizing which a class may be assigned to each of the extracted terms. Specifically, the term-document binding table 200 is shown to include three columns, namely a "number of occurrences" column 202, a density column 204, and an assigned class column 206. A term having a density value of greater than four percent, for example, is identified as falling in the "A" class, a term having a density of between two and four percent is identified as falling in the "B" class, a term having a density of between one and two percent is identified as falling in the "C" class, while a term having a density of between 0.5 and one percent is identified as falling in the "D" class. For the terms having a density of above 0.5 percent, the density value is utilized to assign a class. For terms which have a density value less than 0.5 percent, the count value is utilized for this purpose. Specifically, a term having a count value of greater than 3 is assigned to the "E" class, and a term having a count value of between 1 and 3 is assigned to the "F" class. Accordingly, the assigned class is indicative of the "binding strength" with which the term is associated with or coupled to the electronic document under consideration.

At step 186, a characteristic (or qualitative) indicator in the form of a term weight value is determined, based on characteristics qualities of the term such as those represented by the Type and Part of Speech indications discussed above. While this determination may again be made in any number of ways, FIG. 11 shows an exemplary weight table 210, utilizing which a weight value may be assigned to each of the extracted terms. Specifically, the weight table 210 is shown to include four columns, namely a weight number column 212, a type column 214, a length column 216 and a Part of Speech column 218. By identifying an appropriate combination of type, length and Part of Speech indications, an appropriate term weight value is assigned to each term. In the type column 214, a type "P" indication identifies an environment lexicon term, a type "L" indication identifies a universal lexicon term, and a type "U" indication identifies

a term of unknown grammatical structure for a given length. The entries within the length column 216 indicate the number of words included within the term. The entries within the Part of Speech column 218 indicate the parts of speech that the words within a term comprise. The "A" indication identifies the adjectives, the "V" indication identifies a verb, the "N" indication identifies a noun, and the "X" indication identifies an unknown part of speech. By mapping a specific term to an appropriate entry within the weight table 210, an appropriate term weight value, as indicated in the weight column 212, may be assigned to the term.

At step 188, a relevancy quantitative indicator in the form of an adjusted count value for each term, is calculated, this adjusted count value being derived from the binding strength and term weight values calculated at steps 184 and 186. While this determination may again be made in any number of ways, FIG. 12 shows an exemplary occurrence factor table 220, utilizing which an adjusted count value for the relevant term may be determined. The occurrence factor table 220 is shown to include values for various binding strength/term weight value combinations. The adjusted count value is indicative of the importance or relevance of term within a single, given document, and does not consider the importance or relevance of the term in view of any occurrences of the term in other electronic documents that may be associated with a particular user.

At step 190, a determination is made as to whether any adjusted count values exists for the relevant term as a result of the occurrence of the term in previously received and analyzed documents. If so, the adjusted count values for occurrences of the term in all such previous documents are summed.

At step 192, an initial confidence level values for the term is then determined based on the summed adjusted counts and the term weight, as determined above with reference to the weight table 210 shown in FIG. 11. To this end, FIG. 13 illustrates a confidence level table 230, which includes various initial confidence level values for various summed adjusted count/weight value combinations that may have been determined for a term. For example, a term having a summed adjusted count of 0.125, and a weight value of 300, may be allocated an initial confidence level value of 11.5. Following the determination of an initial confidence level value, confidence level values for various terms may be grouped into "classes", which still retain cardinal meaning, but which standardize the confidence levels into a finite number of "confidence bands". FIG. 14 illustrates a modified table 240, derived from the confidence level table 230, wherein the initial confidence levels assigned are either rounded up or rounded down to certain values. By grouping into classes by rounding, applications (like e-mail addressing), can make use of the classes without specific knowledge/dependence on the numerical values. These can then be tuned without impact to the applications. The modified confidence level values included within the table 240 may have significance in a number of applications. For example, users may request that terms with a confidence level of greater than 1000 automatically be published in a "public" portion of their user knowledge profile. Further, e-mail addressees for a particular e-mail may be suggested based on a match between a term in the e-mail and a term within the user knowledge profile having a confidence level value of greater than, merely for example, 600.

The method 154 then terminates at step 194.

In a further embodiment of the present invention, the method 154, illustrated in FIG. 9A, may be supplemented by

a number of additional steps 195, as illustrated in FIG. 9B, by which a "document weight" value is assigned to a document based on addressee information associated with the document. The document weight value may be utilized in any one of the steps 182-192 illustrated in FIG. 9A, for example, as a multiplying factor to calculate a confidence level value for a term. In one exemplary embodiment, the binding strength value, as determined at step 184, may be multiplied by the document weight value. In another exemplary embodiment, the term weight value, as determined at step 186, may be multiplied by the document weight value.

The document weight value may be calculated by the confidence logic 45 within the term extractor 46. Referring to FIG. 9B, at step 196, the confidence logic 45 identifies the actual addressee information. To this end, the term extractor 46 may include a header parser (not shown) that extracts and identifies the relevant addressee information. At step 197, the confidence logic 45 then accesses a directory structure that may be maintained by an external communication program for the purposes of determining the level of seniority within an organization of the addressees associated with the document. In one exemplary embodiment of the invention, the directory structure may be a Lightweight Directory Access Protocol (LDAP) directory maintained by a groupware server, such as Microsoft Exchange or Lotus Notes. At step 198, a cumulative seniority level for the various addressees is determined by summing seniority values for each of the addressees. At step 199, the summed seniority value is scaled to generate the document weight value. In this embodiment, the cumulative or summed seniority level of the various addressees comprises an "average" seniority value that is used for the purpose of calculating the document weight term. Alternatively, instead of summing in the seniority values at step 198, a "peak" seniority value (i.e., a seniority value based on the seniority level of the most senior addressee) may be identified and scaled at step 199 to generate the document weight value.

In alternative embodiments, the addressee information may be utilized in a different manner to generate a document weight value. Specifically, a document weight value may be calculated based on the number of addressees, with a higher number of addressees resulting in a greater document weight value. Similarly, a document weight value may be calculated based on the number of addressees who are included within a specific organizational boundary (e.g., a specific department or division). For example, an e-mail message addressed primarily to an executive group may be assigned a greater document weight value than an e-mail message addressed primarily to a group of subordinates. Further, the document weight value may also be calculated using any combination of the above discussed addressee information characteristics. For example, the document weight value could be calculated using both addressee seniority and addressee number information.

CONSTRUCTION OF A USER KNOWLEDGE PROFILE

FIG. 15A is a flow chart illustrating a method 250, according to one exemplary embodiment of the present invention, of constructing a user profile that includes first and second portions that may conveniently be identified as "private" and "public" portions. Specifically, unrestricted access to the "public" portion of the user knowledge profile may be provided to other users, while restricted access to the "private" portion may be facilitated. For example, unrestricted access may encompass allowing a user to review details concerning a user knowledge profile, and the target

user, responsive to a specific request and without specific authorization from the target user. Restricted access, on the other hand, may require specific authorization by the target user for the provision of information concerning the user knowledge profile, and the target user, in response to a specific request. The method 250 commences at step 252, and then proceeds to step 254, where a determination is made regarding the confidence level value assigned to a term, for example using the method 154 described above with reference to FIG. 9A. Having determined the confidence level value, the method 250 proceeds to step 256, where a threshold value is determined. The threshold value may either be a default value, or a user specified value, and is utilized to categorize the relevant term. For example, users may set the threshold through the browser interface as a fundamental configuration for their profile. If set low, the user profile will be aggressively published to the public side. If set high, only terms with a high level of confidence will be published. Users can also elect to bypass the threshold publishing concept altogether, manually reviewing each term that crosses the threshold (via the notification manager) and then deciding whether to publish. At decision box 258, a determination is made as to whether the confidence level value for the term is less than the threshold value. If so, this may be indicative of a degree of uncertainty regarding the term as being an accurate descriptor of a user's knowledge. Accordingly, at step 260, the relevant term is then stored in the "private" portion of the user knowledge profile. Alternatively, should the confidence level value be greater than the threshold value, this may be indicative of a greater degree of certainty regarding the term as an accurate descriptor of a user's knowledge, and the relevant term is then stored in the "public" portion of the user's knowledge profile at step 262. The method 150 then terminates at step 264.

FIG. 16A shows an exemplary user-term table 112, constructed according to the method 250 illustrated in FIG. 15A. Specifically, the table 112 is shown to include a first user knowledge profile 270 and a second user knowledge profile 280. The first user knowledge profile 270 is shown to include a "public" portion 272, and a "private" portion 274, the terms within the "private" portion 274 having an assigned confidence level value (as indicated in the confidence level column 118) below a threshold value of 300. The second user knowledge profile 280 similarly has a "public" portion 282 and a "private" portion 284.

The exemplary user-term table 112 shown in FIG. 16A comprises an embodiment of the table 112 in which the public and private portions are determined dynamically with reference to a confidence level value assigned to a particular user-term pairing. FIG. 16B illustrates an alternative embodiment of the user-term table 112 that includes a "private flag" column 119, within which a user-term pairing may be identified as being either public or private, and accordingly part of either the public or private portion of a specific user profile. While the state of a private flag associated with a particular user-term pairing may be determined exclusively by the confidence level associated with the pairing, in an alternative embodiment of the invention, the state of this flag may be set by other mechanisms. For example, as described in further detail below with reference to FIG. 17E, a user may be provided with the opportunity manually to modify the private or public designation of a term (i.e., move a term between the public and private portions of a user knowledge profile). A user may be provided with an opportunity to modify the private or public designation of a term in response to a number of events. Merely for example, a user may be prompted to designate a

term as public in response to a "hit" upon a term in the private portion during a query process, such as during an "expert-lookup" query or during an "addressee-lookup" query. When storing the term in the user knowledge profile at either steps 260 or 262, the allocation of the term to the appropriate portion may be made by setting a flag, associated with the term, in the "private flag" column 119 within the user-term table 112, as illustrated in FIG. 16B. For example, a logical "1" entry within the "private flag" column 119 may identify the associated term as being in the "private" portion of the relevant user knowledge profile, while a logical "0" entry within the "private flag" column 119 may identify the associated term as being in the "public" portion of the relevant user knowledge profile.

(87) FIG. 15B illustrates an exemplary method 260/262, according to one embodiment of the present invention, of storing a term in either a public or private portion of a user knowledge profile. Specifically, a respective term is added to a notification list at step 1264, following the determination made at decision box 258, as illustrated in FIG. 15A. At decision box 1268, a determination is made as to whether a predetermined number of terms have been accumulated within the notification list, or whether a predetermined time period has passed. If these conditions are not met, the method waits for additional terms to be added to the notification list, or for further time to pass, at step 1266, before looping back to the step 1264. On the other hand, should a condition within the decision box 1268 have been met, the method proceeds to step 1270, where the notification list, that includes a predetermined number of terms that are to be added to the user knowledge profile, is displayed to a user. The notification list may be provided to the user in the form of an e-mail message, or alternatively the user may be directed to a web site (e.g., by a URL included within e-mail message) that displays the notification list. In yet a further embodiment, the notification list may be displayed on a web or intranet page that is frequently accessed by the user, such as a home page. At step 1272, the user then selects terms that are to be included in the public portion of the user knowledge profile. For example, the user may select appropriate buttons displayed alongside the various terms within the notification list to identify terms for either the public or private portions of the user knowledge profile. At step 1274, private flags, such as those contained within the "private flag" column 119 of the user-term table 112 as shown in FIG. 16B, may be set to a logical zero "0" to indicate that the terms selected by the user are included within the public portion. Similarly, private flags may be set to a logical one "1" to indicate terms that were not selected by the user for inclusion within the public portion are by default included within the private portion. It will of course be appreciated that the user may, at step 1272, select terms to be included within the private portion, in which case un-selected terms will by default be included within the public portion. The method then ends at step 1280.

The above described method is advantageous in that a user is not required to remember routinely to update his or her user profile, but is instead periodically notified of terms that are candidates for inclusion within his or her user knowledge profile. Upon notification, the user may then select terms for inclusion within the respective public and private portions of the user knowledge profile. As such, the method may be viewed as a "push" model for profile maintenance.

METHOD OF ACCESSING A USER KNOWLEDGE PROFILE

(90) While the above method 250 is described as being executed at the time of construction of a user knowledge

profile, it will readily be appreciated that the method may be dynamically implemented as required and in response to a specific query, with a view to determining whether at least a portion of a user knowledge profile should be published, or remain private responsive to the relevant query. To this end, FIG. 17A shows a flow chart illustrating a method 300, according to one exemplary embodiment of the present invention, of facilitating access to a user knowledge profile. The method 300 commences at step 302, and then proceeds to step 304, where a threshold value is determined. At step 306, a document term within an electronic document generated by a user (hereinafter referred to as a "query" user) is identified. Step 306 is performed by the term extractor 46 responsive, for example, to the receipt of an e-mail from the mail system interface 42 within the knowledge gathering system 28. At step 308, comparison logic 51 within the term extractor 46 identifies a knowledge term within the repository 50 corresponding to the document term identified at step 306. The comparison logic 51 also determines a confidence level value for the identified knowledge term. At decision box 310, the comparison logic 51 makes a determination as to whether the confidence level value for the knowledge term identified at step 308 is less than the threshold value identified at step 304. If not (that is the confidence level value is greater than the threshold value) then a public profile process is executed at step 312. Alternatively, a private profile process is executed at step 314 if the confidence level value falls below the threshold value. The method 300 then terminates at step 316.

FIG. 17B shows a flowchart illustrating an alternative method 301, according to an exemplary embodiment of the present invention, of facilitating access to a user knowledge profile. The method 301 commences at step 302, and then proceeds to step 306, where a document term within an electronic document generated by a user (i.e., the "query" user) is identified. The term extractor 46 performs step 306 responsive, for example, to the receipt of an e-mail message from the mail system interface 42 within the knowledge gathering system 28. At step 308, the comparison logic 51 within the term extractor 46 identifies a knowledge term within the knowledge repository 50 corresponding to the document term identified at step 306. At decision box 311, the comparison logic 51 then makes a determination as to whether a "private" flag for the knowledge term is set to indicate the relevant knowledge term as being either in the public or the private portion of a user knowledge profile. Specifically, the comparison logic 51 may examine the content of an entry in the private flag column 112 of a user-term table for a specific user-term pairing of which the knowledge term is a component. If the "private" flag for the knowledge term is set, thus indicating the knowledge term as being in the private portion of a user knowledge profile, the private profile process is executed at step 314. Alternatively, the public profile process is executed at step 312. The method 301 then terminates at step 316.

FIG. 17C shows a flow chart detailing a method 312, according to an exemplary embodiment of the present invention, of performing the public profile process mentioned in FIGS. 17A and 17B. The method 312 commences at step 320, and user information, the knowledge term corresponding to the document term, and the confidence level value assigned to the relevant knowledge term are retrieved at steps 322, 324, and 326. This information is then displayed to the query user at step 328, whereafter the method 312 terminates at step 330.

(94) FIG. 17D shows a flow chart detailing a method 314, according to an exemplary embodiment of the present

invention, of performing the private profile process mentioned in FIGS. 17A and 17B. The method 314 commences at step 340, and proceeds to step 342, where a user (herein after referred to as the "target" user) who is the owner of the knowledge profile against which the hit occurred is notified of the query hit. This notification may occur in any one of a number of ways, such as for example via an e-mail message. Such an e-mail message may further include a URL pointing to a network location at which further information regarding the query hit, as well as a number of target user options, may be presented. At step 346, the reasons for the query hit are displayed to the target user. Such reasons may include, for example, matching, or similar, document and knowledge terms utilizing which the hit was identified and the confidence level value associated with the knowledge term. These reasons may furthermore be presented within the e-mail propagated at step 342, or at the network location identified by the URL embedded within the e-mail. At step 348, the target user then exercises a number of target user options. For example, the target user may elect to reject the hit, accept the hit, and/or modify his or her user knowledge profile in light of the hit. Specifically, the target user may wish to "move" certain terms between the public and private portions of the user knowledge profile. Further, the user may optionally delete certain terms from the user knowledge profile in order to avoid any further occurrences of hits on such terms. These target user options may furthermore be exercised via a HTML document at the network location identified by the URL. At decision box 350, a determination is made as to whether the user elected to modify the user knowledge profile. If so, a profile modification process, which is described below with reference to FIG. 17E, is executed at step 352. Otherwise, a determination is made at decision box 354 as to whether the target user rejected the hit. If so, the hit is de-registered at step 356. Alternatively, if the target user accepted the hit, the public profile process described above with reference to FIG. 17C is executed at step 358. The method 314 then terminates at step 360.

FIG. 17E is a flowchart illustrating a method 352, according to an exemplary embodiment of the present invention, for implementing the profile modification process illustrated at step 352 in FIG. 17D. The method 352 commences at step 362, and then proceeds to display step 364, where the target user is prompted to (1) move a term, on which a "hit" has occurred, between the private and public portions of his or her user knowledge profile, or to (2) delete the relevant term from his or her user knowledge profile. Specifically, the target user may be presented with a user dialog, a HTML-enriched e-mail message, or a Web page, listing the various terms upon which hits occurred as a result of an inquiry, besides which appropriate buttons are displayed that allow the user to designate the term either to the included in the public or private portion of his or her user knowledge profile, or that allow the user to mark the relevant term for deletion from the user knowledge profile. At input step 366, the target user makes selections regarding the terms in the matter described above. At decision box 368, a determination is made as to whether the user selected terms for transfer between the public and private portions of the user profile, or for inclusion within the user profile. If so, the method 352 proceeds to step 370, wherein the appropriate terms are designated as being either public or private, in accordance with the user selection, by setting appropriate values in the "private flag" column 119 within the user-term table, as illustrated in FIG. 16B. Thereafter, the method proceeds to decision box 372, wherein a determination is made as to

whether the user has elected to delete any of the terms presented at step 364. If so, the relevant terms are deleted from the user knowledge profile at step 374. The method is then terminates at step 378.

The methodologies described above with reference to FIGS. 15 through 17E are advantageous in that, where the confidence level of a term falls below a predetermined threshold, the owner of the user knowledge profile may elect to be involved in the process of determining whether a query hit is accurate or inaccurate. The owner of the user knowledge profile is also afforded the opportunity to update and modify his or her knowledge profile as and when needed. Further, the owner of the user knowledge profile is only engaged in the process for hits below a predetermined certainty level and on a public portion of the knowledge profile. Matches between document terms and knowledge terms in the public portion are automatically processed, without any manual involvement.

METHOD FOR ADDRESSING AN ELECTRONIC DOCUMENT FOR TRANSMISSION OVER A NETWORK

Returning now briefly to FIG. 5, the addressing system 84 within the e-mail client extensions 19 operates independently of the profiling system 82 to suggest potential recipients for an e-mail message based on the content thereof. The user interface 80 within the e-mail client extensions 19 may pop-up a window when the system determines such suggestion is possible, based on the length of a draft message being sent, or may present a command button labeled "Suggest Recipients". This button is user selectable to initiate a sequence of operations whereby the author of the e-mail is presented with a list of potential recipients who may be interested in receiving the e-mail based on predetermined criteria, such as a match between the content of the e-mail and a user profile, or a commonality with a confirmed addressee.

FIG. 18A is a flow chart illustrating a method 400, according to an exemplary embodiment of the present invention, of addressing an electronic document, such as an e-mail, for transmission over a network, such as the Internet or an Intranet. The method 400 commences at step 402, and then proceeds to step 401, where a determination is made as to whether the body of the draft message exceeds a predetermined length (or number of words). If so, content of the electronic document (e.g., an e-mail message body) is transmitted to the knowledge access server 26 via the web server 20 at step 404. Specifically, a socket connection is open between the e-mail client 18 and the web server 20, and the content of the message body, which may still be in draft form, is transmitted using the Hypertext Transfer Protocol (HTTP) via the web server 20 to the knowledge access server 26. At step 406, the knowledge access server 26 processes the message body, as will be described in further detail below. At step 408, the knowledge access server 26 transmits a potential or proposed recipient list and associated information to the addressing system 84 of the e-mail client 18. Specifically, the information transmitted to the e-mail client 18 may include the following:

1. A list of user names, as listed within column 94 of the user table 90, as well as corresponding e-mail addresses, as listed within the column 98 of the user table 90;
2. A list of term identifiers, as listed in column 116 of the user-term table 112, that were located within the "public" portion of a user knowledge profile that formed the

basis for a match between document terms within the message body and knowledge terms within the user knowledge profile; and

3. A "matching metric" for each user included in the list of user names (1). Each "matching metric" comprises the sum of the confidence level values, each multiplied by the weighted occurrences of the term within the message body, for the terms identified by the list of term identifiers (2) and associated with the relevant user. This "matching metric" is indicative of the strength of the recommendation by the knowledge access server 26 that the relevant user (i.e., potential recipient) be included within the list of confirmed addressees.

At step 410, the author of the electronic document is presented with a list of potential recipients by the e-mail client 18, and specifically by the addressing system 84 via a user dialog 440 as shown in FIG. 18D. FIG. 18D groups matching levels into matching classes each characterized by a visual representation (icon).

The user dialog 440 shown in FIG. 18D presents the list of potential recipients in a "potential recipients" scrolling window 442, wherein the names of potential recipients are grouped into levels or ranked classes according to the strength of the matching metric. An icon is also associated with each user name, and provides an indication of the strength of the recommendation of the relevant potential recipients. Merely for example, a fully shaded circle may indicate a high recommendation, with various degrees of "blackening" or darkening of a circle indicating lesser degrees of recommendation. A "rejection" icon may be associated with an actual recipient, and an example of such a "rejection" icon is indicated at 441. The "rejection" icon indicates a negative recommendation on an actual recipient supplied by the author of the message, and may be provided in response to a user manually modifying his or her profile to designate certain terms therein as generating such a "rejection" status for a recipient against which a hit occurs.

The user dialog 440 also presents a list of actual (or confirmed) recipients in three windows, namely a "to:" window 442, a "cc:" window 444 and a "bcc:" window 446. An inquiring user may move recipients between the potential recipients list and the actual recipients lists utilizing the "Add" and "Remove" buttons indicated at 450. The user dialog 440 also includes an array of "select" buttons 452, utilizing which a user can determine the recommendation group to be displayed within the scrolling window 442. The user dialog 440 finally also includes "Explained Match" and "More" buttons 454 and 456, the purposes of which is elaborated upon below. As shown in FIG. 18D, the author user may select an "Explain" function for any of the proposed recipients utilizing the "Explain Match" button 454. If it is determined at decision box 412 that this "Explain" function has been selected, the method 400 branches to step 414, as illustrated in FIG. 18B. Specifically, at step 414, the addressing system 84 propagates a further "Explain" query to the knowledge access server 26 utilizing HTTP, and opens a browser window within which to display the results of the query. At step 416, the knowledge access server 26 retrieves the terms (i.e., the knowledge terms) that constituted the basis for the match, as well as associated confidence level values. This information is retrieved from the public portion of the relevant user knowledge profile in the knowledge repository 50. At step 418, the information retrieved at step 416 is propagated to the client 18 from the knowledge access server 26 via the web server 20. The information is then displayed within the browser window opened by the e-mail

client 18 at step 414. Accordingly, the author user is thus able to ascertain the reason for the proposal of a potential recipient by the addressing system 84, and to make a more informed decision as to whether the proposed recipient should be included within the actual recipients (confirmed addressee) list.

The user also has the option of initiating a "More" function by selecting the "More" button 456 on the user dialog 440, this function serving to provide the user with additional proposed recipients. Accordingly, a determination is made at step 422 as to whether the "More" function has been selected by the author user. If so, the method 400 branches to step 424 as shown in FIG. 18C, where the client 18 propagates a "More" request to the knowledge access server 20 in the same manner as the "Explain" query was propagated to the knowledge access server at step 414. At step 426, the knowledge access server 26 identifies further potential recipients, for example, by using a threshold value for the "matching metric" that is lower than a threshold value utilized as a cutoff during the initial information retrieval operation performed at steps 406 and 408. At step 428, the knowledge access server 26 then transmits the list of further potential recipients, and associated information, to the e-mail client 18. At step 430, the list of additional potential recipients is presented to the author user for selection in descending order according to the "matching metric" associated with each of the potential recipients.

At step 432, the user then adds at his or her option, or deletes selected potential or "rejected" recipients to the list of actual recipients identified in "to:", "cc:" or "bcc:" lists of the e-mail, thus altering the status of the potential recipients to actual recipients. At step 434, the e-mail message is then transmitted to the confirmed addressees.

If the user profile includes a "rejection" status on a term (something a user can do through manual modification of the profile), then a special symbol, such as that indicated 441 in FIG. 18D, may be returned indicating a negative recommendation on a recipient supplied by the author of the message.

The exemplary method 400 discussed above is advantageous in that the knowledge access server 26 automatically provides the author user with a list of potential addressees, based on a matching between document terms identified within the message body of an e-mail and knowledge terms included within user profiles.

CASE CONTROL

FIG. 19 is a flow chart illustrating a method 500, according to one exemplary embodiment of the present invention, of managing user authorization to publish, or permit access to, a user knowledge profile. The method 500 is executed by the case controller 45A that tracks open "cases" and initiates notification to users concerning the status of such cases. For the purposes of the present specification, the term "case" may be taken to refer to a user authorization process for publication of, or access to, a user knowledge profile. The method 500 commences at step 502, and then proceeds to step 504, where a match is detected with a private portion of a user knowledge profile. At step 504, the case controller 45A then opens a case, and notifies the target user at step 506 concerning the "hits" or matches between a document (or query) term and a knowledge term in a knowledge user profile. This notification may be by way of an e-mail message, or by way of publication of information on a Web page accessed by the user. At step 508, the case controller 45A determines whether an expiration date, by which the target user is required to respond to the hit, has been reached

or in fact passed. If the expiration date has passed, the case controller 45A closes the case and the method 500 terminates. Alternatively, a determination is made at decision box 510 as to whether the target user has responded to the notification by authorizing publication of, or access to, his or her user knowledge profile based on the hit on the private portion thereof. If the target user has not authorized such action (i.e., declined authorization), an inquiring user (e.g., the author user of an e-mail or a user performing a manual database search to locate an expert) is notified of the decline at step 512. Alternatively, should the target user have authorized publication or access, the inquiring user is similarly notified of the authorization at step 514. The notification of the inquiring user at steps 512 or 514 may be performed by transmitting an e-mail to the inquiring user, or by providing a suitable indication on a web page (e.g., a home page or search/query web page) accessed by the inquiring user. At step 516, the appropriate portions of the user profile pertaining to the target user are published to the inquiring user, or the inquiring user is otherwise permitted access to the user profile. At step 518, the case controller 45A then closes the case, whereafter the method terminates.

SUPPLEMENTAL METHOD OF IDENTIFYING CONFIDENCE VALUE

FIGS. 7-9 describe an exemplary method 140 of identifying knowledge terms and calculating associated confidence level values. A supplemental method 550, according to an exemplary embodiment of the present invention, of assigning a confidence value to a term will now be described with reference to FIGS. 20-22. The supplemental method 550 seeks to compensate for a low confidence level value which may be associated with the term as a result of the term not appearing in any recent documents associated with a user. It will be appreciated that by calculating a confidence level value utilizing the method illustrated in FIG. 9, aged terms (i.e., terms which have not appeared in recent documents) may be attributed a low confidence level value even though they may be highly descriptive of a specialization or knowledge of a user. The situation may occur where a user is particularly active with respect to a particular topic for a short period of time, and then re-focuses attention on another topic. Over time, the methodology illustrated in FIG. 9 may too rapidly lower the confidence level values associated with terms indicating user knowledge.

Referring to FIG. 20, there is illustrated the exemplary method 550 of assigning a confidence value to a term. The method 550 commences at step 552, whereafter an initial confidence memory value (as distinct from a confidence level value) is assigned a zero (0) value. At step 556, a confidence level value for a term is calculated utilizing, for example, the method 154 illustrated in FIG. 9. However, this confidence level value is only calculated for occurrences of the relevant term within a particular time or document window. For example, in summing the adjusted count values at step 190 within the method 154, the adjusted count values for only documents received within a predetermined time (e.g., the past 30 days), or only for a predetermined number of documents (e.g., the last 30 documents) are utilized to calculate the summed adjusted count value. It will be appreciated that by discarding documents, which occurred before the time or document window, the effect on the confidence level values for aged terms by the absence of such aged terms within recent documents may be reduced.

At decision box 558, a determination is then made as to whether a newly calculated confidence level value for a term is greater than a previously recorded confidence memory

value, or alternatively greater than a predetermined site-wide or system-wide threshold value. If the confidence level value is determined to be greater than the confidence memory value (or the threshold value), the confidence memory value is then made equal to the confidence level value by overwriting the previous confidence memory value with the newly calculated confidence level value. In this way, it is ensured that the confidence level value does not exceed the confidence memory value.

FIG. 22 is an exemplary user-term table 112, according to one embodiment of the present invention, that is shown to include a confidence level column 118, a confidence memory value column 121, and a time stamp column 123. The table 122 records a confidence level value and a confidence memory value for each user-term pairing within the table 112, and it is to this table that the confidence level values and the confidence memory values are written by the method 550. The time stamp column 123 records a date and time stamp value indicative of the date and time at which the corresponding confidence memory value was last updated. This value will accordingly be updated upon the overwriting of the confidence memory value at step 560.

Should the confidence level value not exceed the confidence memory value or the threshold value, as determined at decision box 558, the method 550 then proceeds to decision box 562, where a further determination is made as to whether another time or document window, associated with a step of decaying the confidence memory value, has expired. If not, the confidence memory value is left unchanged at step 564. Alternatively, if the time or document window associated with the decay step has expired, the confidence memory value is decayed by a predetermined value or percentage at step 566. For example, the confidence memory value may be decayed by five (5) percent per month. The time stamp value may be utilized to determine the window associated with the decay step. The time stamp value associated with the decayed confidence memory value is also updated at step 566. The method 550 then terminates at step 568.

FIG. 21 is a flowchart illustrating an exemplary method 570, according to one embodiment of the present invention, of determining or identifying a confidence value (e.g., either a confidence level value or a confidence memory value) for a term. The method 570 may be executed in performance of any of the steps described in the preceding flow charts that require the identification of a confidence level value for a term in response to a hit on the term by a document term (e.g., in an electronic document or other query). The method 570 commences at step 572, and proceeds to step 574, where a confidence level value for a term within a user profile is identified. For example, the confidence level value may be identified within the user-term table 112 illustrated in FIG. 22. At step 576, a confidence memory value for the term may then also be identified, again by referencing the user-term table 112 illustrated in FIG. 22. At decision box 578, a determination is then made as to whether the confidence level value is greater than the confidence memory value. If the confidence level value is greater than the confidence memory value, the confidence level value is returned, at step 580, as the confidence value. Alternatively, should the confidence memory value be greater than the confidence level value, the confidence memory value is returned, at step 582, as the confidence value. The method 570 then terminates at step 584.

Accordingly, by controlling the rate at which a confidence value for a term is lowered or decayed, the present invention seeks to prevent having a potentially relevant term ignored or overlooked.

COMPUTER SYSTEM

FIG. 23 is a diagrammatic representation of a machine in the form of computer system 600 within which software, in the form of a series of machine-readable instructions, for performing any one of the methods discussed above may be executed. The computer system 600 includes a processor 602, a main memory 603 and a static memory 604, which communicate via a bus 606. The computer system 600 is further shown to include a video display unit 608 (e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT)). The computer system 600 also includes an alphanumeric input device 610 (e.g., a keyboard), a cursor control device 612 (e.g., a mouse), a disk drive unit 614, a signal generation device 616 (e.g., a speaker) and a network interface device 618. The disk drive unit 614 accommodates a machine-readable medium 615 on which software 620 embodying any one of the methods described above is stored. The software 620 is shown to also reside, completely or at least partially, within the main memory 603 and/or within the processor 602. The software 620 may furthermore be transmitted or received by the network interface device 618. For the purposes of the present specification, the term "machine-readable medium" shall be taken to include any medium that is capable of storing or encoding a sequence of instructions for execution by a machine, such as the computer system 600, and that causes the machine to performing the methods of the present invention. The term "machine-readable medium" shall be taken to include, but not be limited to, solid-state memories, optical and magnetic disks, and carrier wave signals.

Thus, a method and apparatus for constructing and maintaining a user knowledge profile have been described. Although the present invention has been described with reference to specific exemplary embodiments, it will be evident that various modifications and changes may be made to these embodiments without departing from the broader spirit and scope of the invention. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. A computer-implemented method of constructing a user knowledge profile, the method including:

automatically assigning a confidence level to content within an electronic document associated with a first user, the content being potentially indicative of a user knowledge base of the first user; and

storing the content in either a first or a second portion of a user knowledge profile of the first user according to the assigned confidence level,

wherein the first and the second portions of the user knowledge profile of the first user have different access restrictions with respect to a second user.

2. The method of claim 1 wherein access to the first portion of the user knowledge profile by the second user is unrestricted and the storing comprises storing the content in the first portion if the confidence level exceeds a predetermined threshold.

3. The method of claim 2 wherein access to the second portion of the user knowledge profile by the second user is restricted and the storing comprises storing the content in the second portion if the confidence level is below the predetermined threshold.

4. The method of claim 1 including designating the first portion of the user knowledge profile of the first user as a public portion and designating the second portion of the user knowledge profile of the first user as a private portion.

5. The method of claim 1 including determining whether the electronic document is indicated by the first user as being for use in construction of the user knowledge profile.

6. The method of claim 5 wherein the determining comprises examining a flag associated with the electronic document, the flag marking the electronic document as being for use in the construction of the user knowledge profile of the first user.

7. The method of claim 5 wherein the determining comprises determining whether the electronic document is addressed to an entity that performs the construction of the user knowledge profile of the first user.

8. The method of claim 1 including determining whether the electronic document indicates an attachment as being for use in the construction of the user knowledge profile of the first user.

9. The method of claim 1 wherein automatically assigning the confidence level comprises identifying a term within the content.

10. The method of claim 9 wherein automatically assigning the confidence level includes counting a number of words within the term.

11. The method of claim 9 wherein automatically assigning the confidence level includes identifying parts of speech comprising the term.

12. The method of claim 9 wherein automatically assigning the confidence level includes comparing the term to a collection of lexicon terms.

13. The method of claim 9 wherein automatically assigning the confidence level includes determining a frequency with which the term occurs within the electronic document associated with the first user.

14. The method of claim 13 wherein the determining of the frequency comprises determining a percentage value indicating a frequency of occurrence of the term relative to a total number of terms within the electronic document associated with the first user.

15. The method of claim 9 wherein automatically assigning the confidence level includes determining a frequency with which the term occurs within a plurality of electronic documents associated with the first user.

16. The method of claim 9 wherein the term comprises any one of a group consisting of a grammar term, a noun phrase and a word.

17. The method of claim 1 wherein automatically assigning the confidence level includes analyzing addressee information associated with the electronic document.

18. The method of claim 17 wherein the analyzing of the addressee information includes identification of a level of seniority of an addressee of the electronic document.

19. The method of claim 17 wherein the analyzing of the addressee information includes determining an average level of seniority of all addressees of the electronic document.

20. The method of claim 17 wherein the analyzing of the addressee information includes determining a level of seniority of a most senior addressee of the electronic document.

21. The method of claim 18 including generating a document weight term utilizing the level of seniority of the addressee of the electronic document.

22. The method of claim 21 including identifying a number of addressees of the electronic document.

23. The method of claim 22 including generating a document weight term utilizing the number of addressees of the electronic document.

24. The method of claim 17 wherein the analyzing of the addressee information includes identification of an organizational group of at least one addressee of the electronic document.

25. The method of claim 23 including generating a document weight term utilizing the organizational group of the at least one addressee of the electronic document.

26. The method of claim 1 wherein the storing of the content includes notifying the first user of the content, prior to storing the content in either the first or second portion of the user knowledge profile of the first user.

27. The method of claim 26 wherein the notifying includes providing the first user with an option of storing the content in either the first or second portion of the user knowledge profile of the first user.

28. The method of claim 26 wherein notifying the first user includes determining whether the confidence level automatically assigned to the content exceeds a predetermined threshold, and notifying the first user of the content when the confidence level exceeds the predetermined threshold.

29. The method of claim 1 wherein the electronic document comprises an electronic mail message generated by the first user.

30. The method of claim 1 wherein the electronic document comprises an electronic document attached to an electronic mail message transmitted by the first user.

31. The method of claim 1 wherein the electronic document comprises a database inquiry generated by the first user.

32. The method of claim 1 wherein the electronic document comprises any one of a group consisting of a user profile document and a list of bookmarks, folders and directories generated by the first user.

33. A data processing system for constructing a user knowledge profile, the system comprising:

confidence logic to examine an electronic document, associated with a first user, and to assign a confidence level to content within the electronic document, the content being potentially indicative of a user knowledge base of the first user; and

a profiler to store the content in either the first or second portion of the user knowledge profile of the first user according to the assigned confidence level,

wherein the first and the second portions of the user knowledge profile of the first user have different access restrictions with respect to a second user.

34. The system of claim 33 wherein access to the first portion of the user knowledge profile by the second user is unrestricted, and the profiler stores the content in the first portion if the confidence level exceeds a predetermined threshold.

35. The system of claim 34 wherein access to the second portion of the user knowledge profile by the second user is restricted, and the profiler stores the content in the second portion if the confidence level is below the predetermined threshold.

36. The system of claim 33 including a client to propagate the electronic document to the confidence logic.

37. The system of claim 33 wherein the client is an e-mail client program.

38. The system of claim 33 wherein the confidence logic determines whether the electronic document is indicated by the first user as being for use in the construction of the user knowledge profile.

39. The system of claim 38 wherein the confidence logic examines a flag associated with the electronic document, the flag marking the electronic document as being for use in the construction of the user knowledge profile of the first user.

40. The system of claim 38 wherein the confidence logic determines whether the electronic document is addressed to

an entity that performs the construction of the user knowledge profile of the first user.

41. The system of claim 33 wherein the confidence logic determines whether the electronic document indicates an attachment as being for use in the construction of the user knowledge profile of the first user.

42. The system of claim 33 wherein the confidence logic identifies a term within the content.

43. The system of claim 42 wherein the confidence logic counts a number of words within the term.

44. The system of claim 42 wherein the confidence logic identifies parts of speech comprising the term.

45. The system of claim 42 wherein the confidence logic compares the term to a collection of lexicon terms.

46. The system of claim 42 wherein the confidence logic determines a frequency with which the term occurs within the electronic document associated with the first user.

47. The system of claim 46 wherein the confidence logic determines a percentage value indicating a frequency of occurrence of the term relative to a total number of terms within the electronic document associated with the first user.

48. The system of claim 42 wherein the confidence logic determines a frequency with which the term occurs within a plurality of electronic documents associated with the first user.

49. The system of claim 42 wherein the term comprises any one of a group consisting of a grammar term, a noun phrase and a word.

50. The system of claim 33 wherein the confidence logic analyzes addressee information associated with the electronic document.

51. The system of claim 50 wherein the confidence logic identifies a level of seniority of an addressee of the electronic document.

52. The system of claim 50 wherein the confidence logic determines an average level of seniority of all addressees of the electronic document.

53. The system of claim 50 wherein the confidence logic determines a level of seniority of a most senior addressee of the electronic document.

54. The system of claim 50 wherein the confidence logic generates a document weight term utilizing the level of seniority of the addressee of the electronic document.

55. The system of claim 33 wherein the confidence logic identifies the number of addressees of the electronic document.

56. The system of claim 55 wherein the confidence logic generates a document weight term utilizing the number of addressees of the electronic document.

57. The system of claim 33 wherein the confidence logic identifies an organizational group of at least one addressee of the electronic document.

58. The system of claim 57 wherein the confidence logic generates a document weight term utilizing the organizational group of the at least one addressee of the electronic document.

59. The system of claim 33 including a notifier that notifies the first user of the content, prior to the content being stored in either the first or the second portion of the user profile of the first user.

60. The system of claim 59 wherein the notifier provides the first user with an option of storing the content in either the first or the second portion of the user knowledge profile of the first user.

61. The system of claim 59 wherein the notifier determines whether the confidence level automatically assigned to the content by the confidence logic exceeds a predeter-

31

mined threshold, and notifies the first user of the content when the confidence level exceeds the predetermined threshold.

62. The system of claim 33 wherein the electronic document comprises an electronic mail generated by the first user.

63. The system of claim 33 wherein the electronic document comprises a document attached to an electronic mail message transmitted by the user.

64. The system of claim 33 wherein the electronic document comprises a database inquiry generated by the first user.

65. The system of claim 33 wherein the electronic document comprises any one of a group consisting of a user profile document and a list of bookmarks, folders, and directories generated by the first user.

66. A data processing system for constructing a user knowledge profile, the system comprising:

confidence means for examining an electronic document, associated with a first user, and for assigning a confidence level to content within the electronic document, the content being potentially indicative of a user knowledge base of the first user; and

32

profiling means for storing the content in either the first or the second portion of the user knowledge profile of the first user according to the assigned confidence level,

wherein the first and the second portions of the user knowledge profile of the first user have different access restrictions with respect to a second user.

67. A computer-readable medium for use by a computer for constructing a user knowledge profile storing a sequence of instructions that, when executed by a computer, cause the computer to perform the steps of:

automatically assigning a confidence level to content within an electronic document associated with a first user, the content being potentially indicative of a user knowledge base of the first user; and

storing the content in either the first or second portion of a user knowledge profile of the first user according to the assigned confidence level,

wherein the first and the second portions of the user knowledge profile of the first user have different access restrictions with respect to a second user.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 6,115,709
DATED : September 5, 2000
INVENTOR(S) : Gilmour et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 29,

Line 38, replace "the" before "first" with -- a --.

Column 32,

Line 1, replace "the" before "first" with -- a --;


Line 2, replace "the" before "second portion" with -- a --;

Line 16, replace "the" before "first or second portion" with -- a --.

Signed and Sealed this

Ninth Day of July, 2002

Attest:

A handwritten signature in black ink, appearing to read "James E. Rogan", written over a horizontal line.

Attesting Officer

JAMES E. ROGAN
Director of the United States Patent and Trademark Office